Matching Supercomputing to Progress in Science

Erik P. DeBenedictis

Sandia National Laboratories SAND2004-3333P

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Introduction

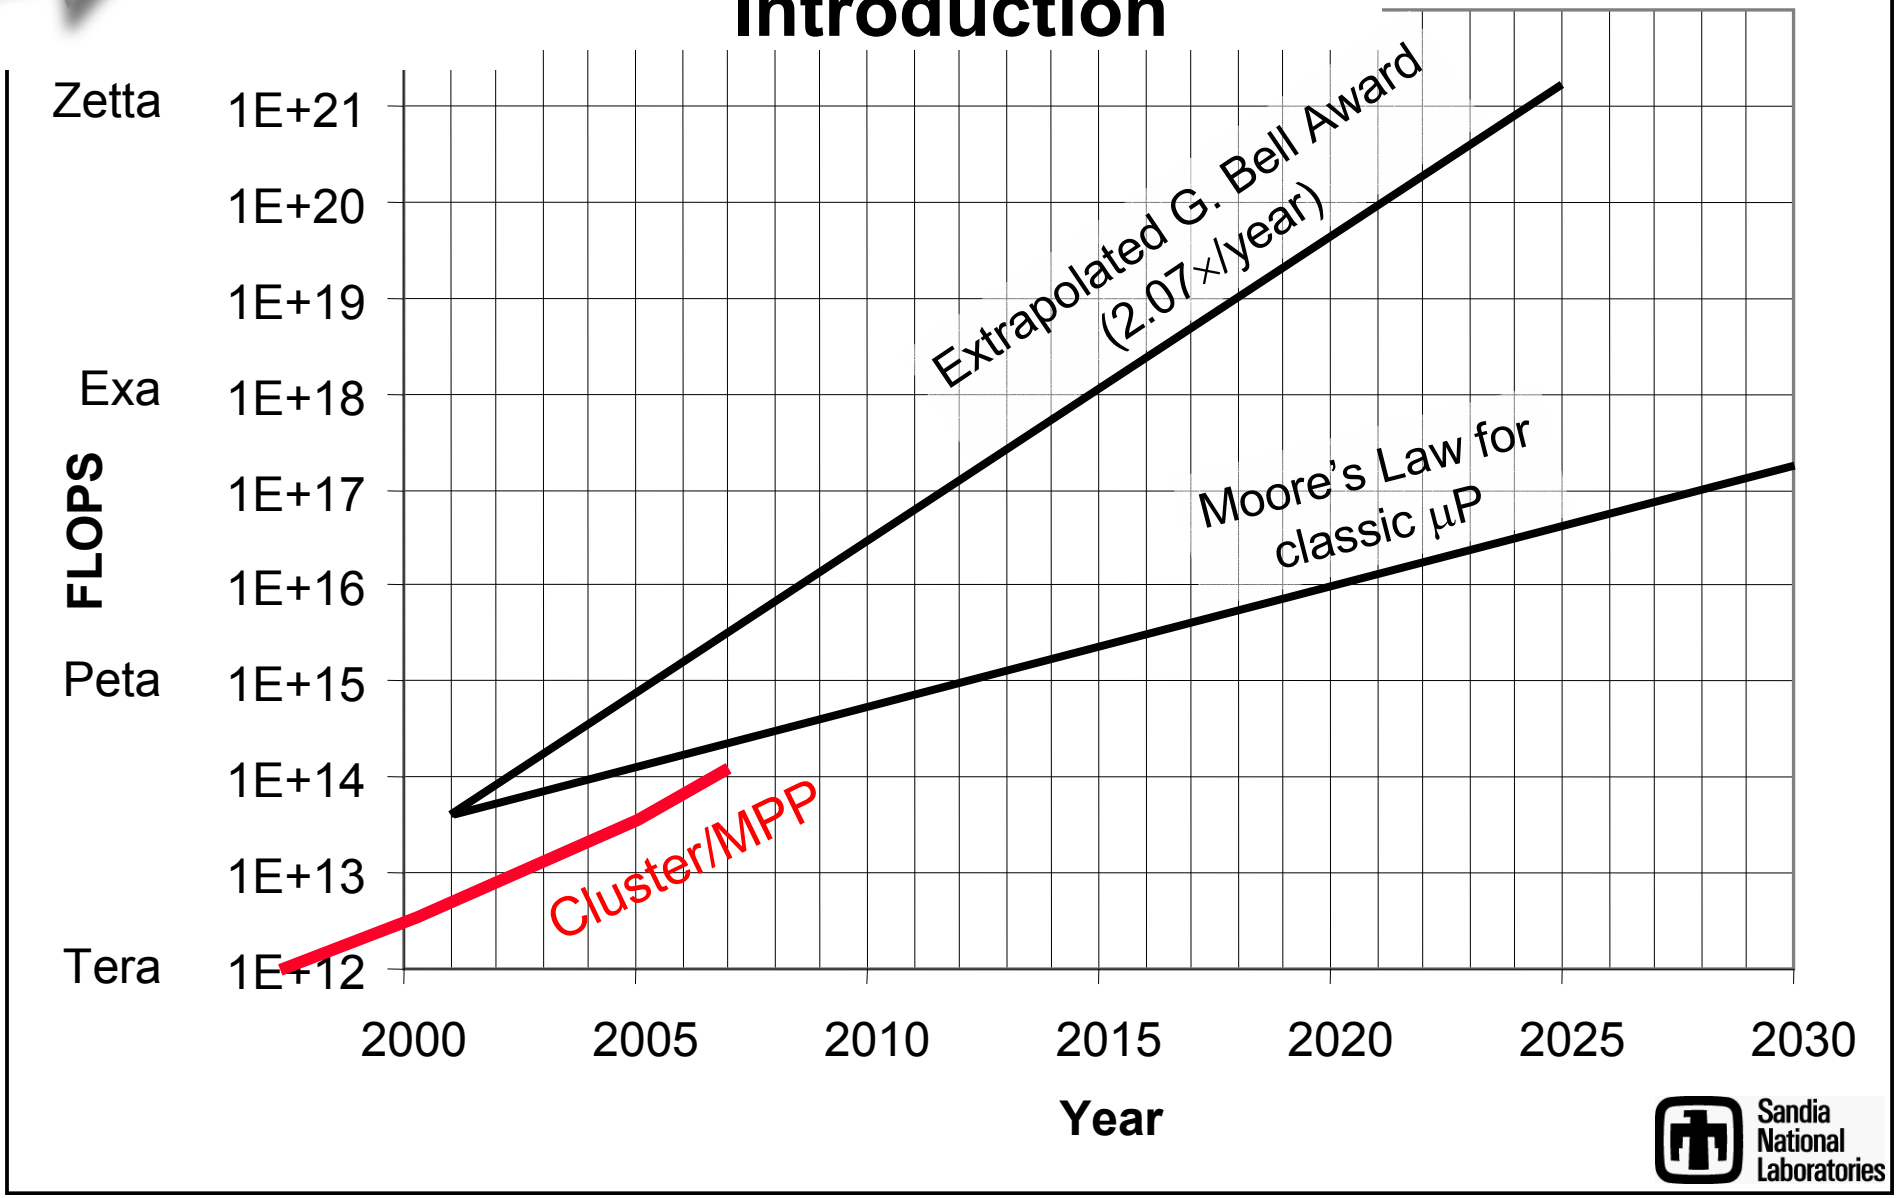

Sandia National Laboratories

Introduction

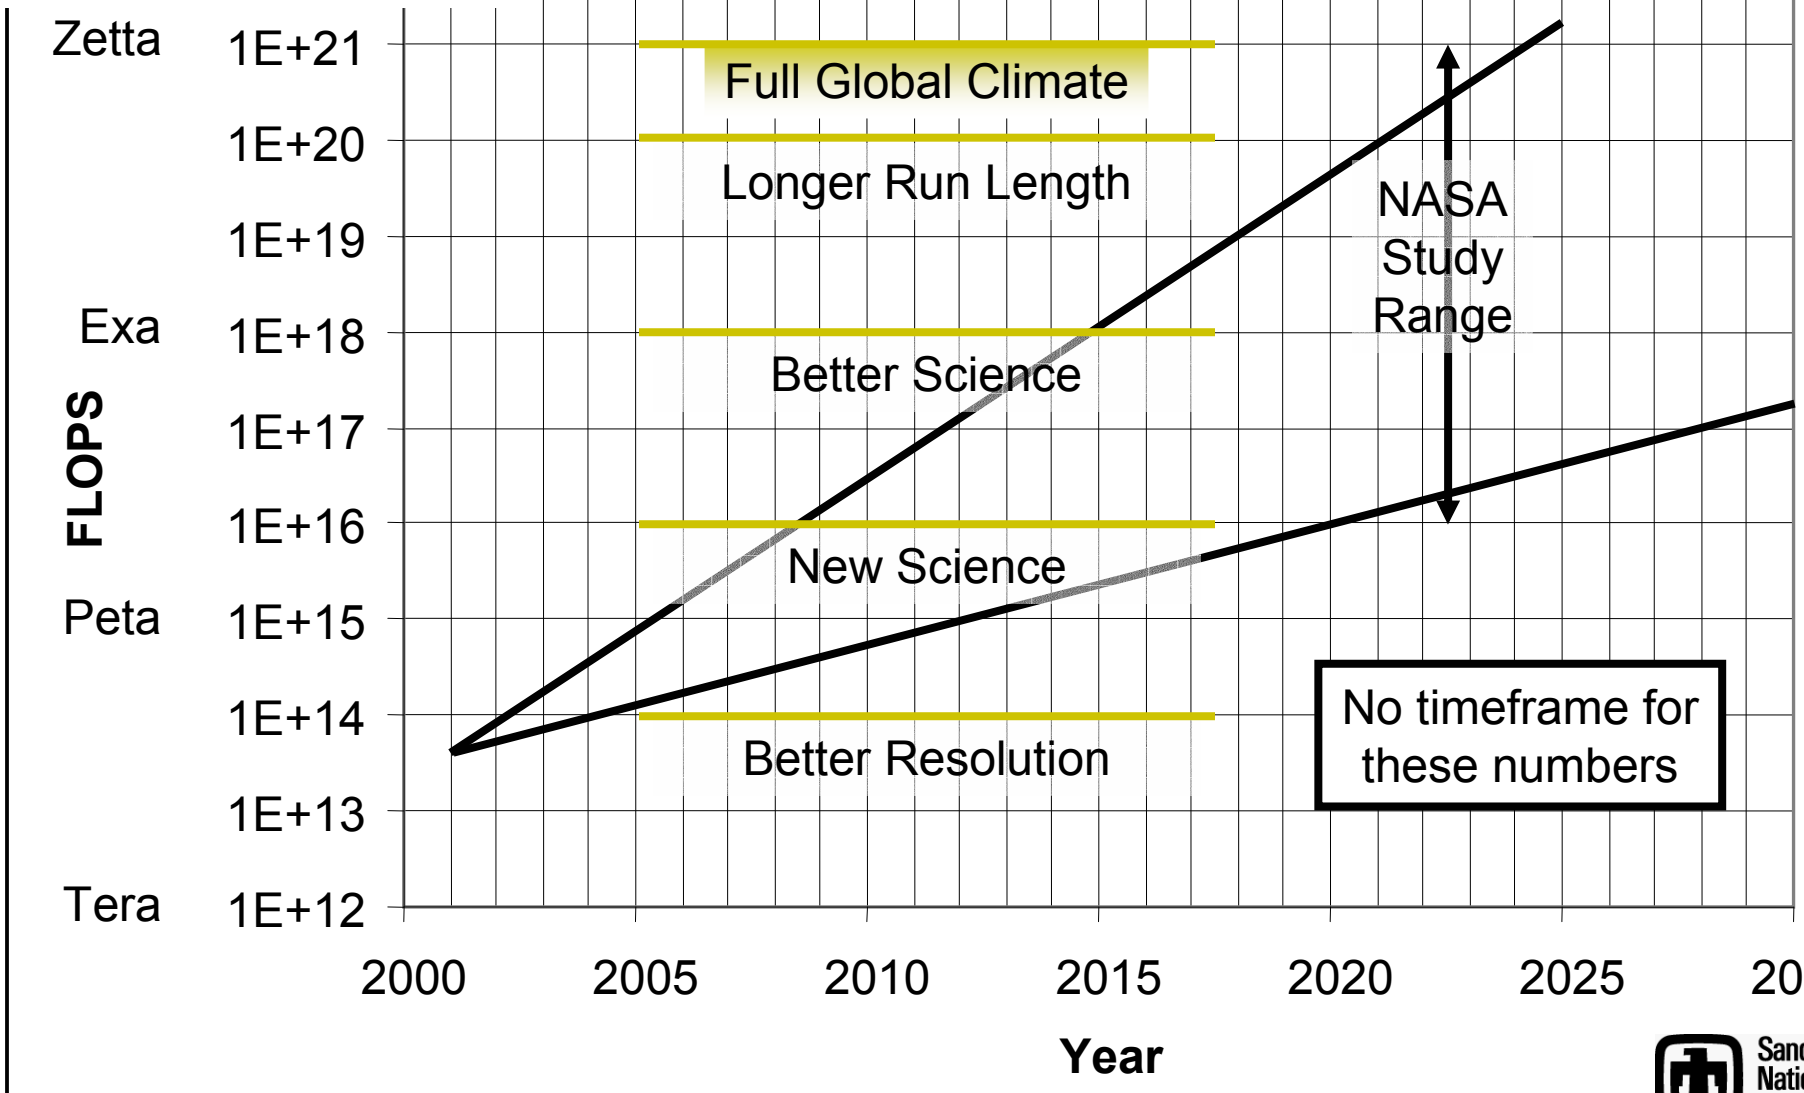

Sandia National Laboratories

Introduction

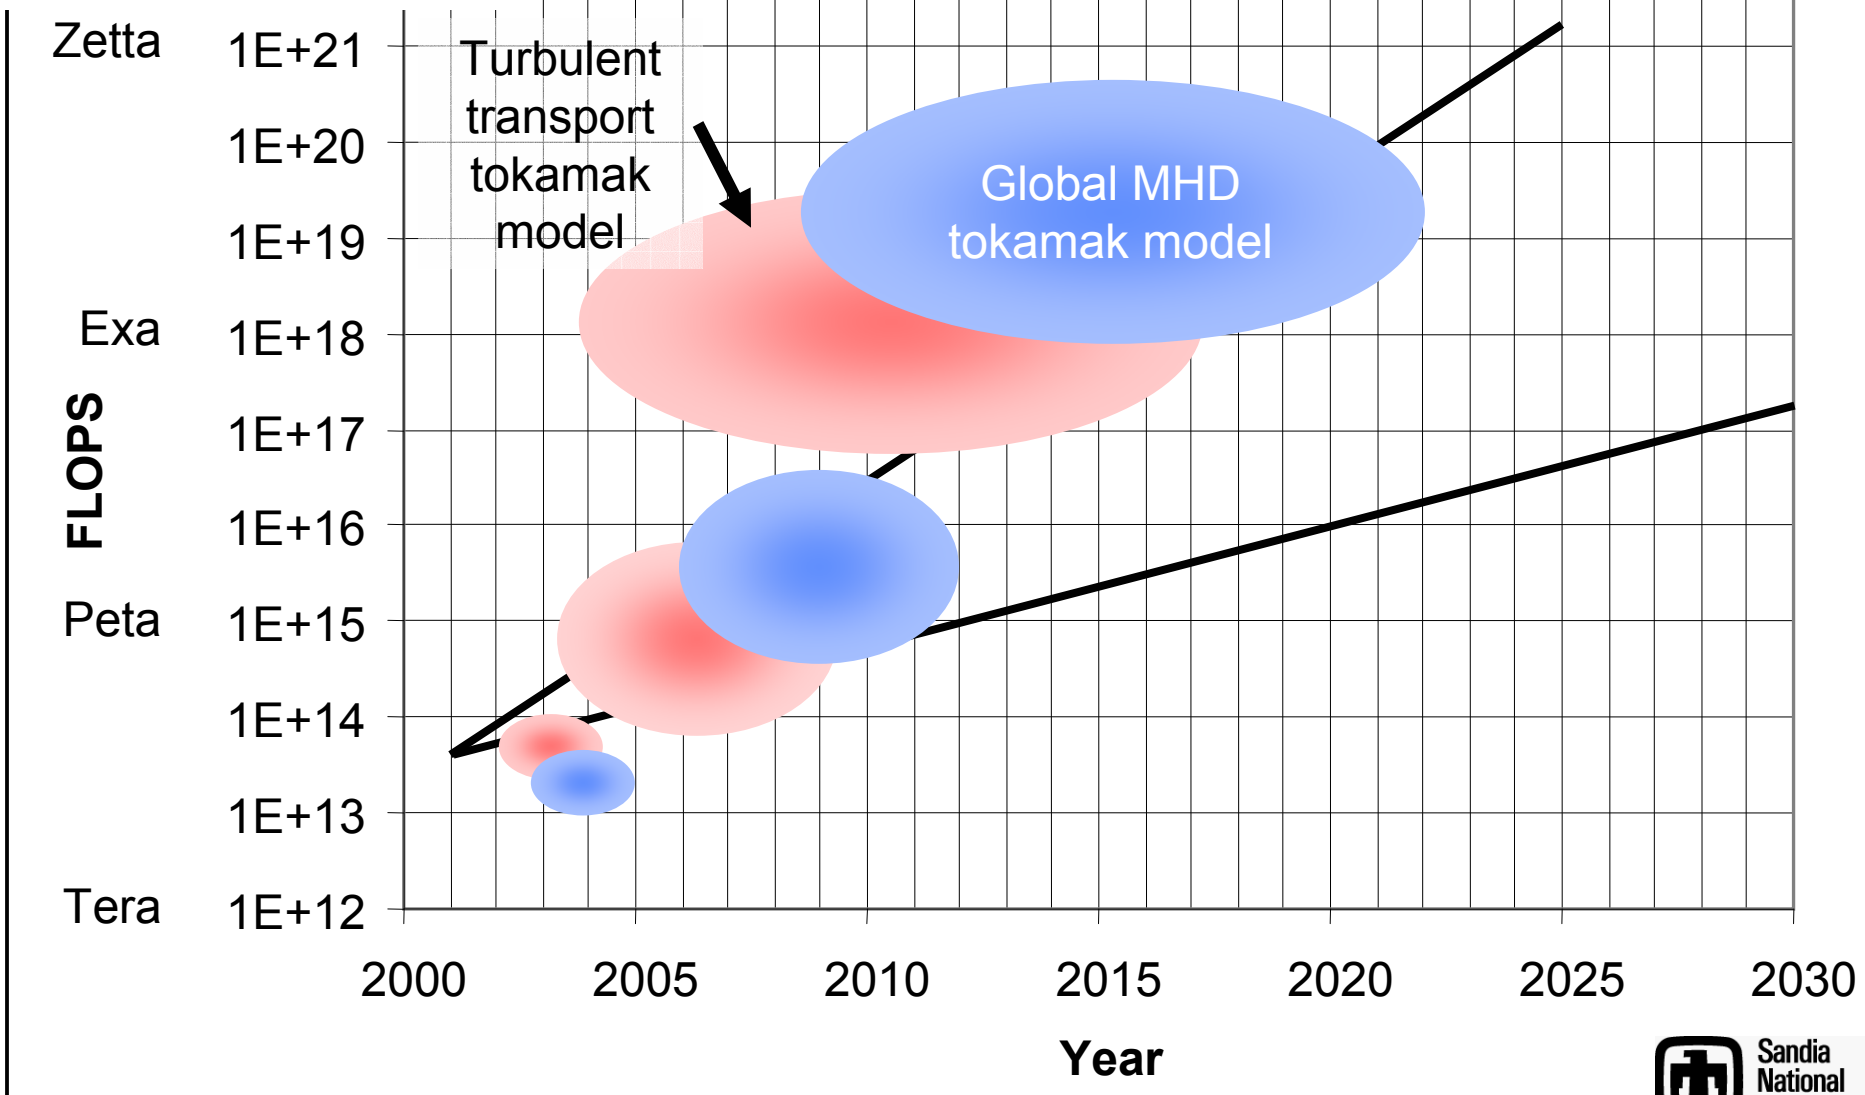

Sandia National Laboratories

Introduction

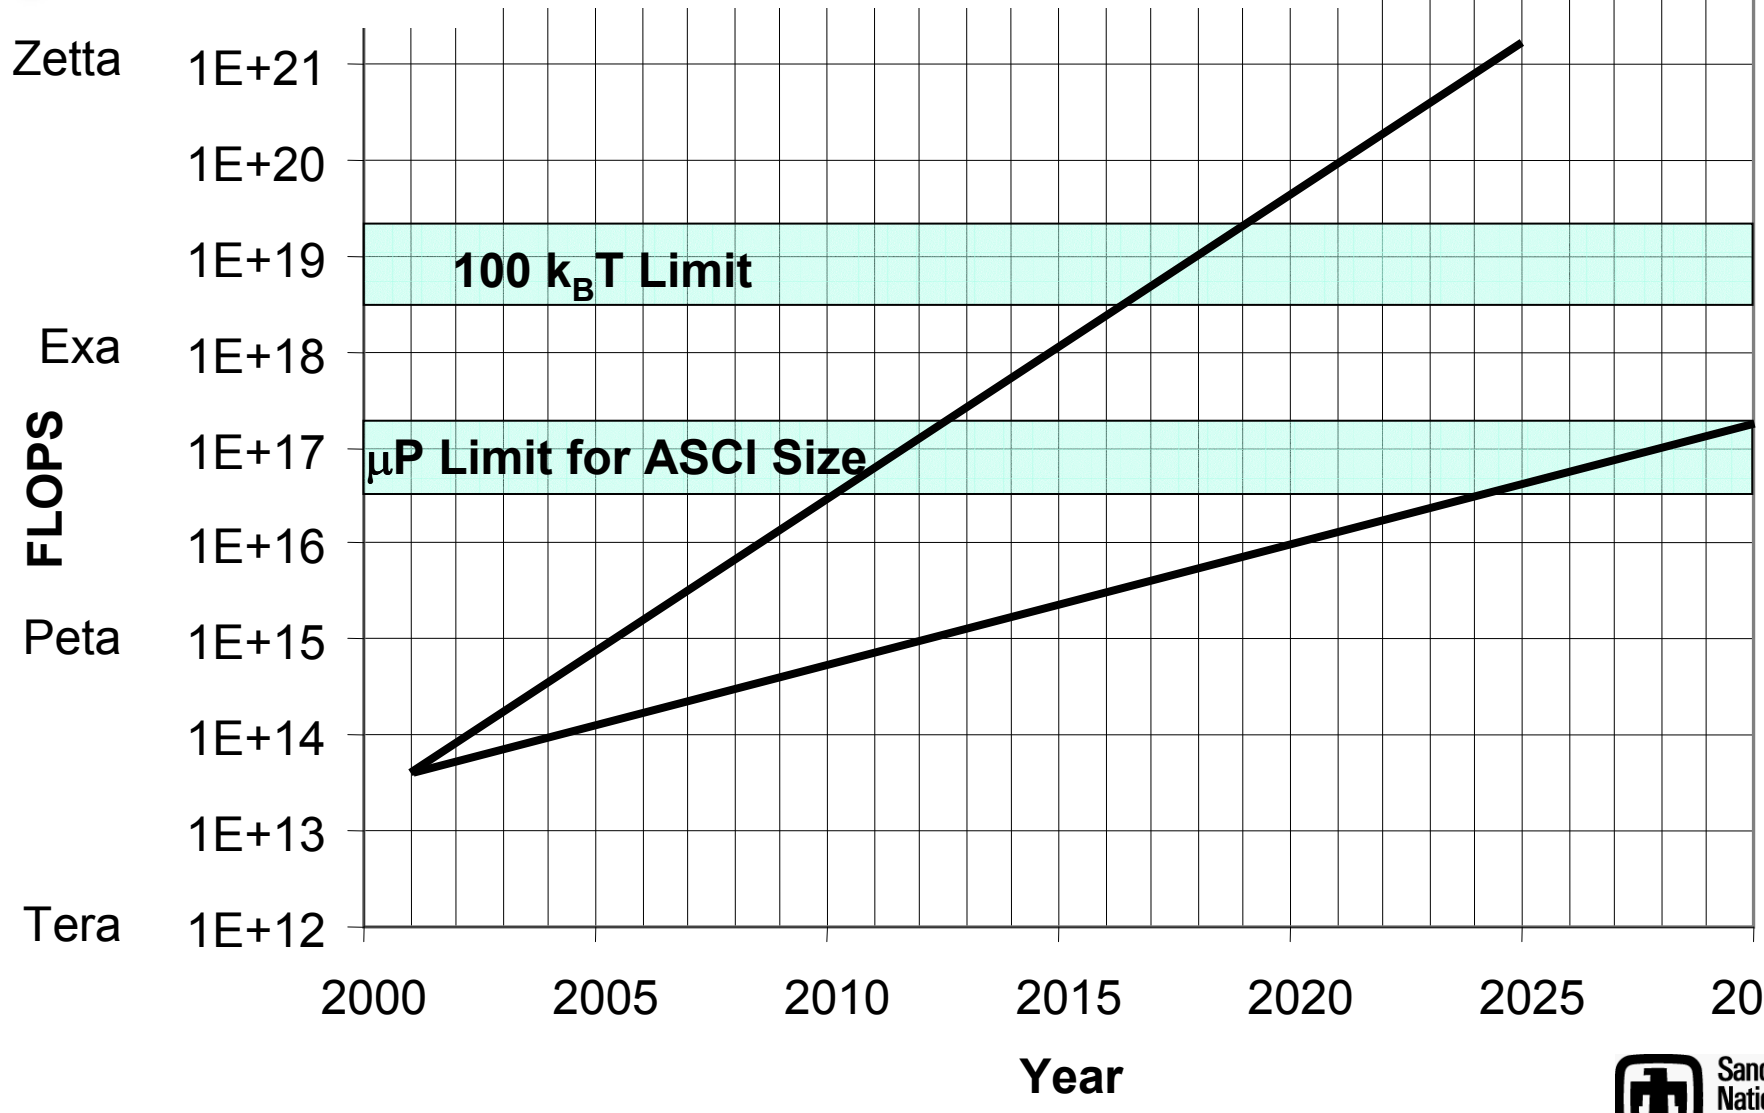

Sandia National Laboratories

Introduction

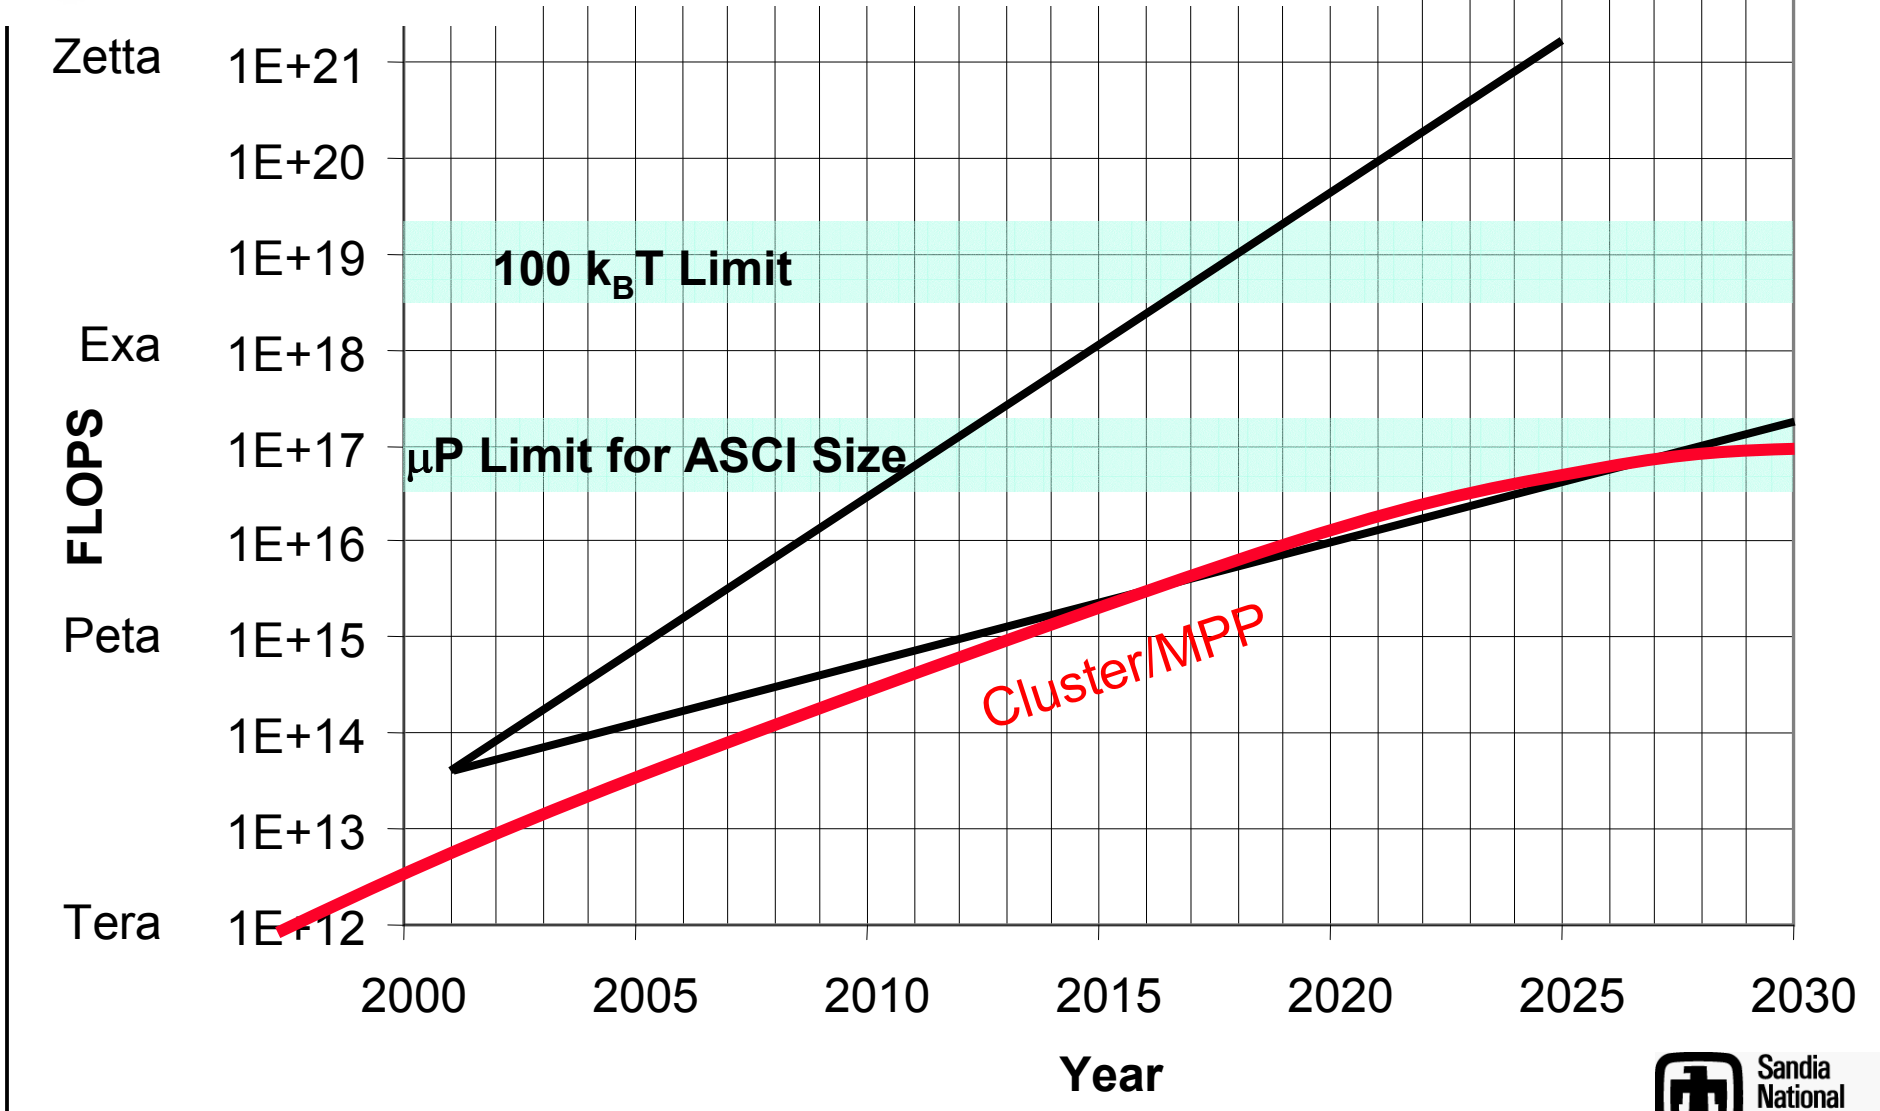

Sandia National Laboratories

Introduction

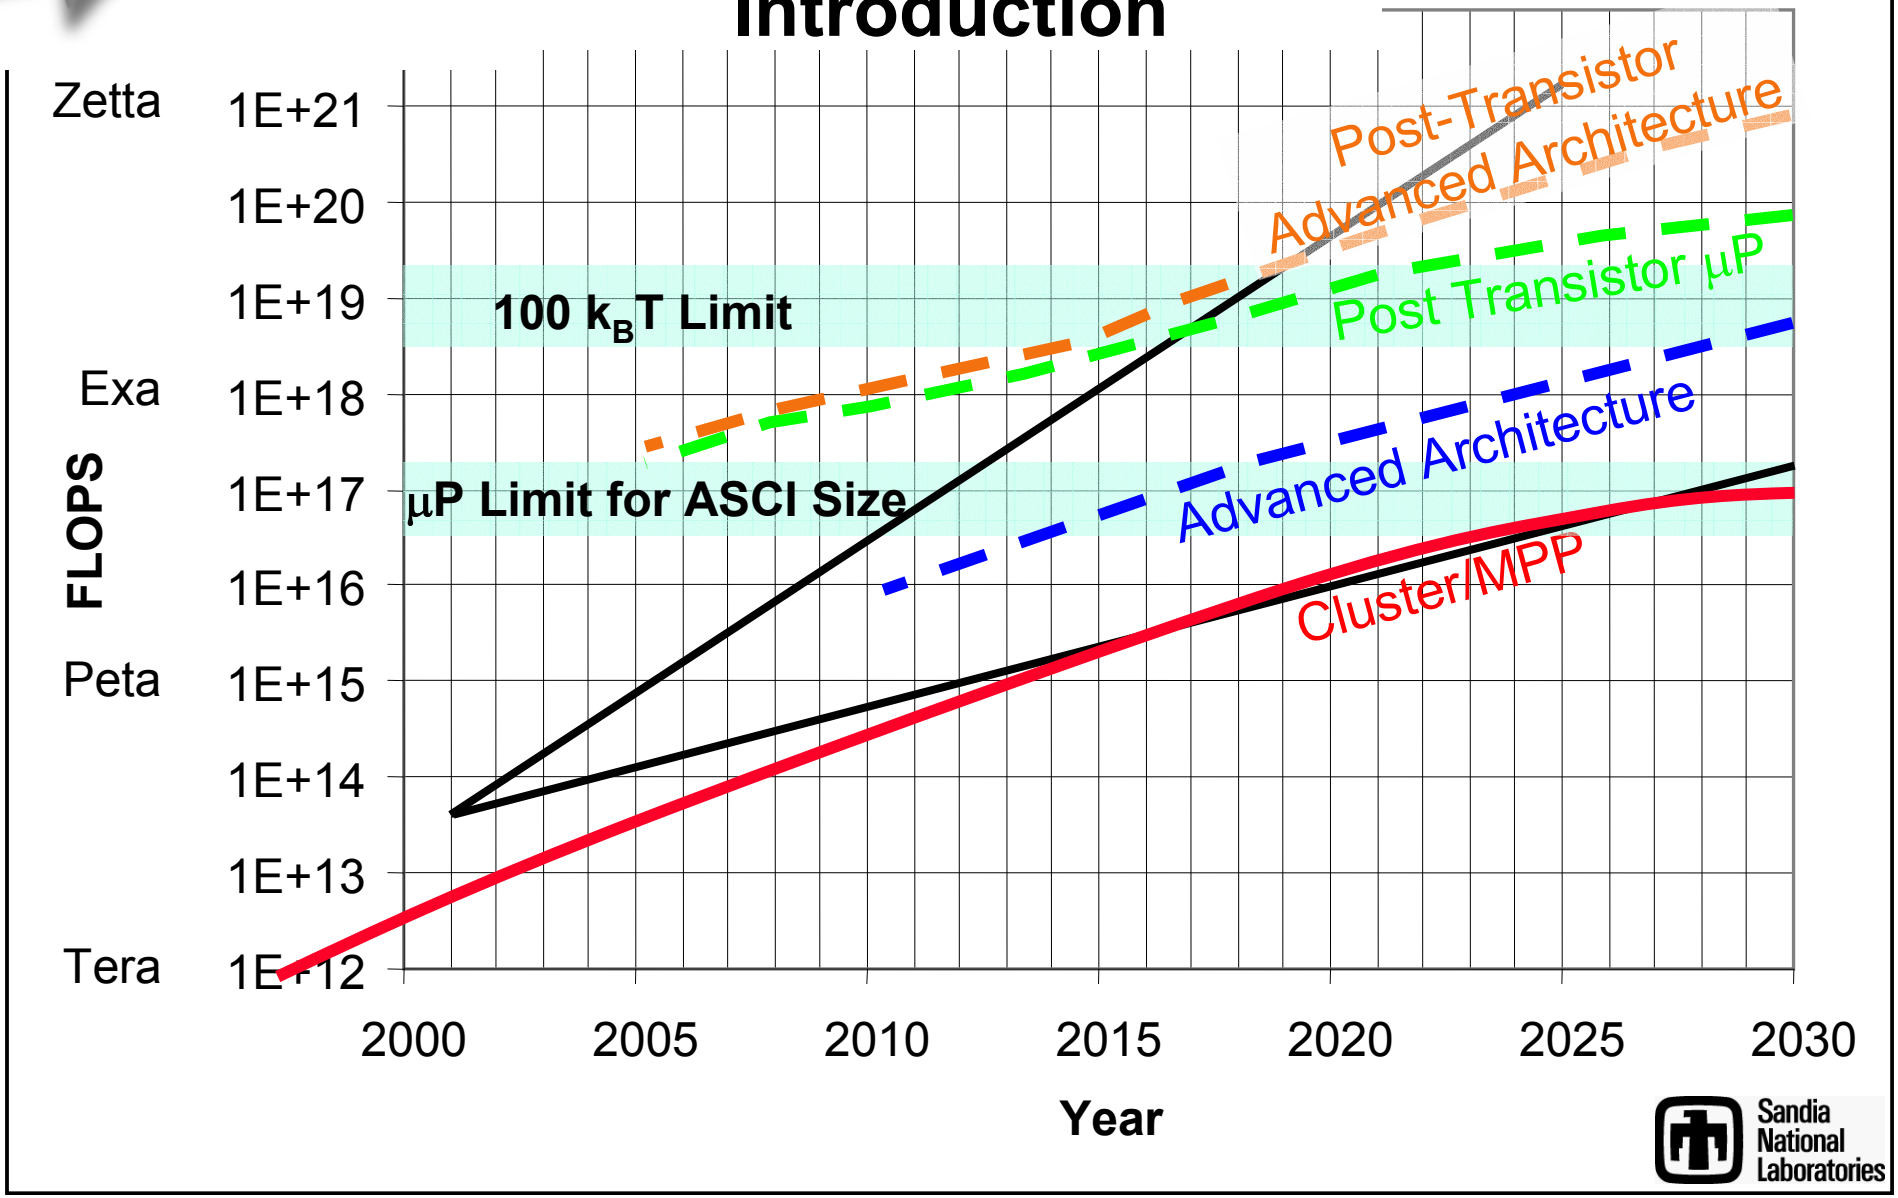

Sandia National Laboratories

Introduction

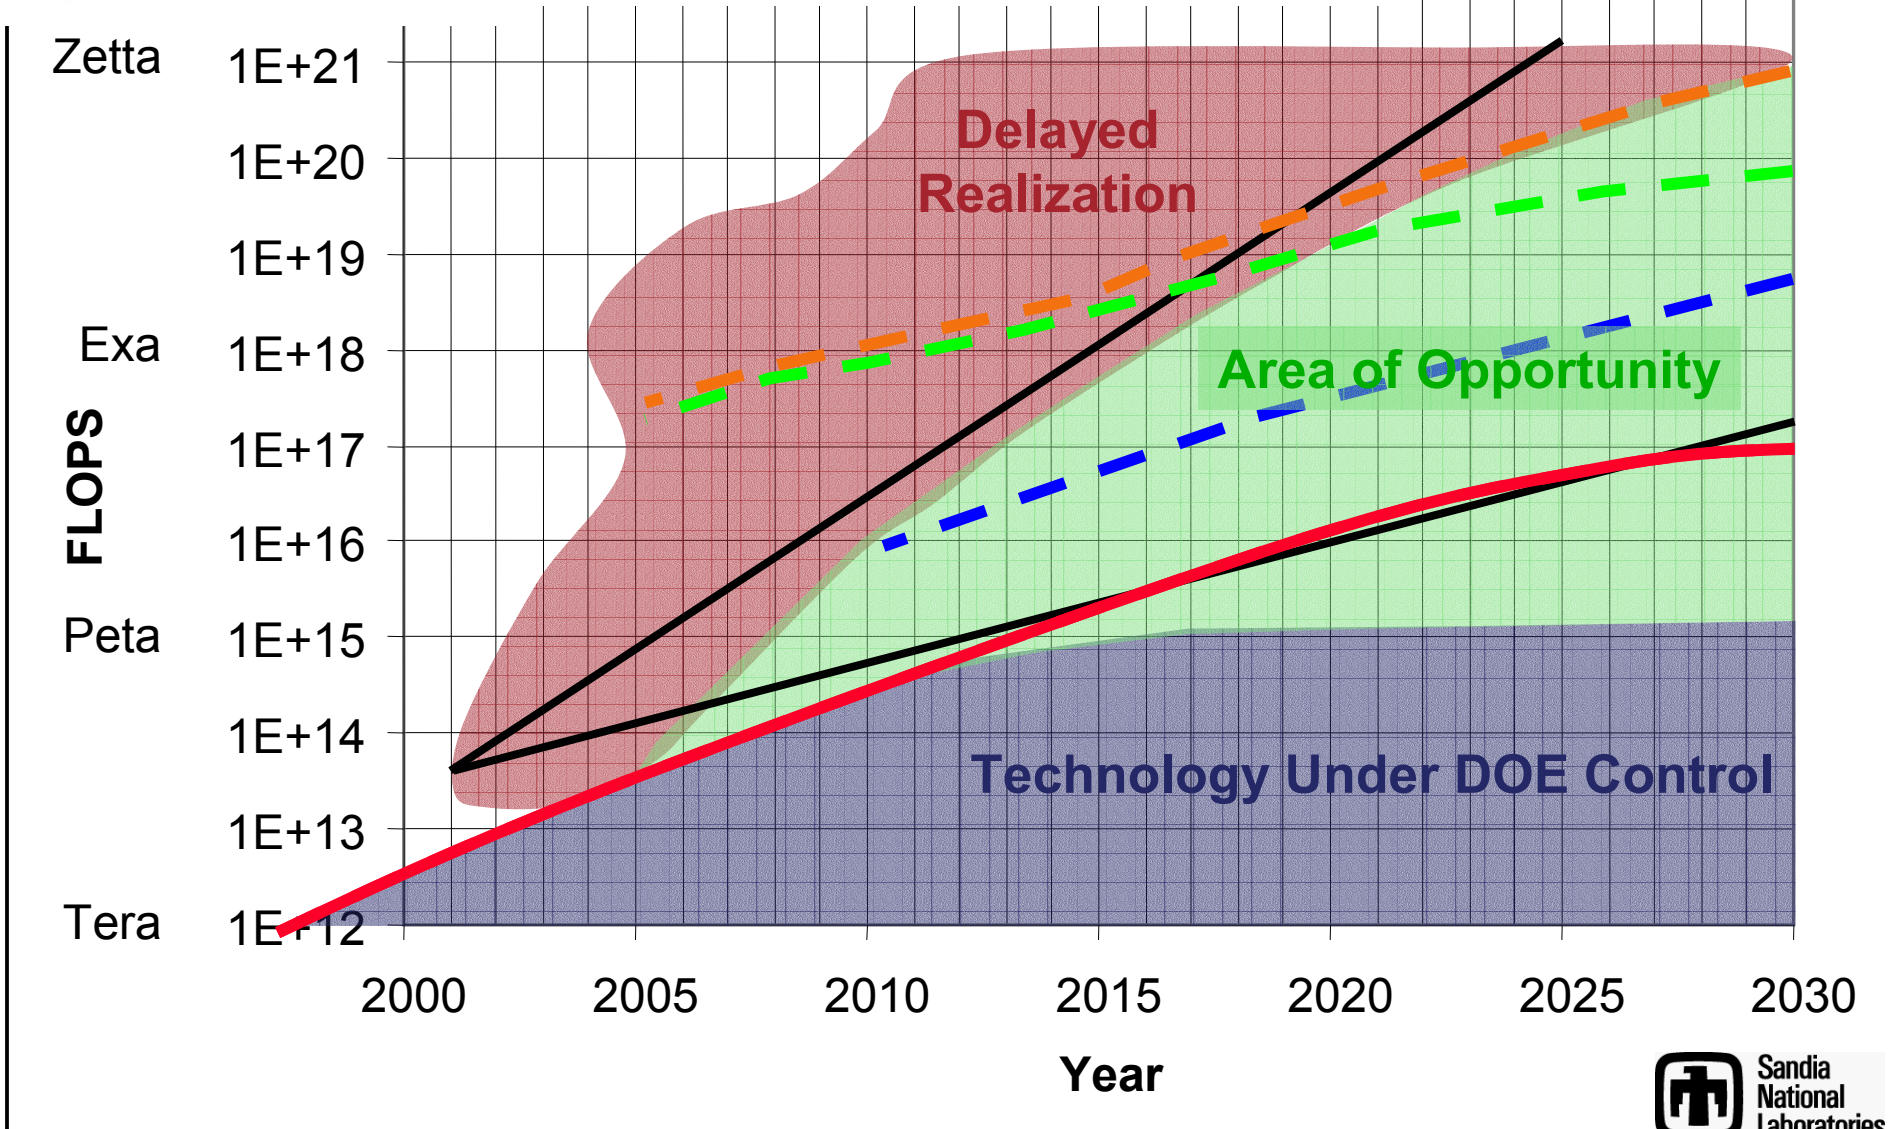

Sandia National Laboratories

Outline

- Applications of the Future
- Limits of Moore's Law
- An Expert System/Optimizer for Supercomputing
- Reaching to Zettaflops
- Roadmap and Future Directions

Sandia National Laboratories

Global Climate

- Objective
  - Collect data about Earth
  - Model climate into the future
  - Provide “decision support” and ability to “mitigate”
- Approaches
  - Climate models exist, but need they more resolution, better physics, and better initial conditions (observations of the Earth)
- Computer Resources Required
  - Increments over current workstation on next slide

FLOPS Increases for Global Climate

	Issue	Scaling
1 Zettaflops	Ensembles, scenarios 10×	Embarrassingly Parallel
100 Exaflops	Run length 100×	Longer Running Time
1 Exaflops	New parameterizations 100×	More Complex Physics
10 Petaflops	Model Completeness 100×	More Complex Physics
100 Teraflops	Spatial Resolution $10^4\times$ ($10^3\times$-$10^5\times$)	Resolution
10 Gigaflops	Current	

Ref. "High-End Computing in Climate Modeling," Robert C. Malone, LANL, John B. Drake, ORNL, Philip W. Jones, LANL, and Douglas A. Rotman, LLNL (2004)

Sandia National Laboratories

Requirements for Plasma Simulation

- Very high peak perform ance requirements
  - but seeking algor- ithmic improvement
- Two methods
  - Red regions very scalable, Monte Carl
  - Green regions N^4 scaling (FEM)
- Long term objective
  - Merge methods into a single code

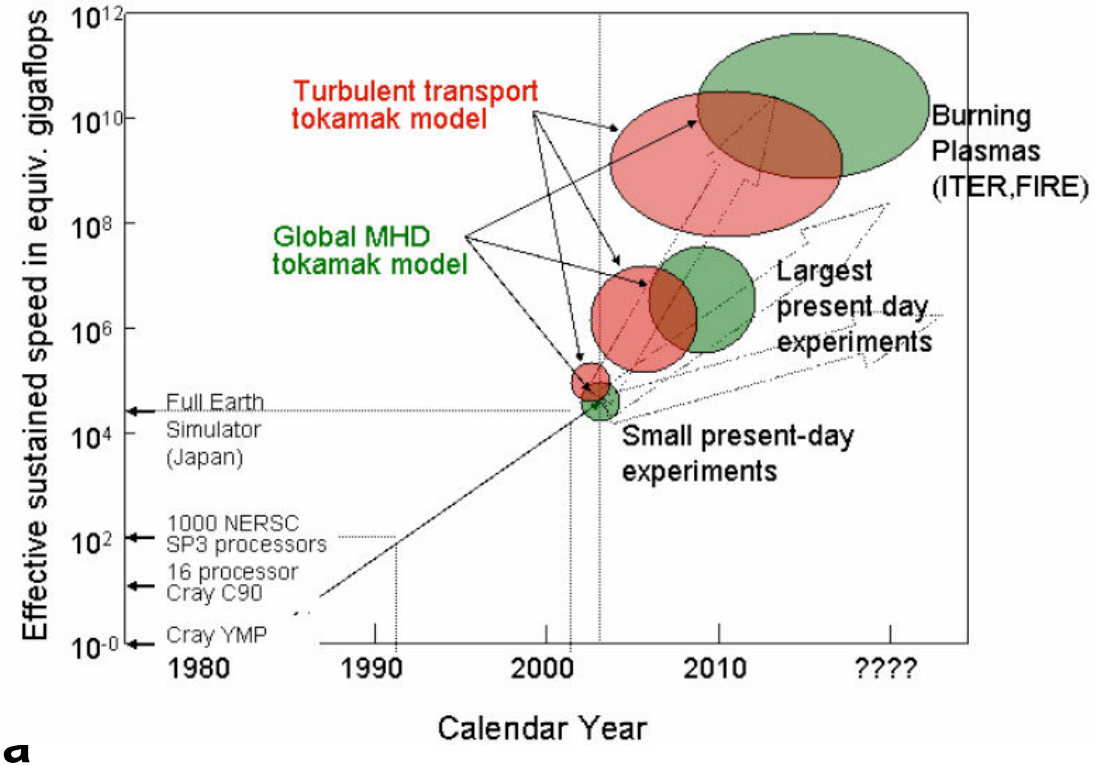

Ref. "Plasma Science Contribution to the SCaLeS Report," S.C. Jardin, October 2003

Sandia National Laboratories

NASA Climate Earth Station

Based on these inputs, various portions of the Modeling and Data Assimilation System will require anywhere from 10^7 to 10^{13} GFLOPS of computational resources. In other words, the range of computational resources needed is 10^{16} to 10^{21} Floating Point Operations per Second. For the curious, the range can also be stated as 10 PetaFLOPS to 1 ZettaFLOPS.

4.1.2. Anticipated Computing Technology Capabilities

At first glance, the numbers discussed in the previous section appear so high as to be impossibly ludicrous. However, with the expected growth in computing capabilities, the lower end of this spectrum actually falls within the domain of possibility.

– **"Advanced Weather Prediction Technologies: NASA's Contribution to the Operational Agencies," Gap Analysis Appendix, May 31, 2002**

Sandia National Laboratories

NASA Work Station

- **"…the ultimate goal of making the computing underlying the design process so capable that it no longer acts as a brake on the flow of the creative human thought…"**
- **Requirement 3 Exaflops**
- **Note: In the context of this report, this requirement is for one or a few engineers, not a supercomputer center!**

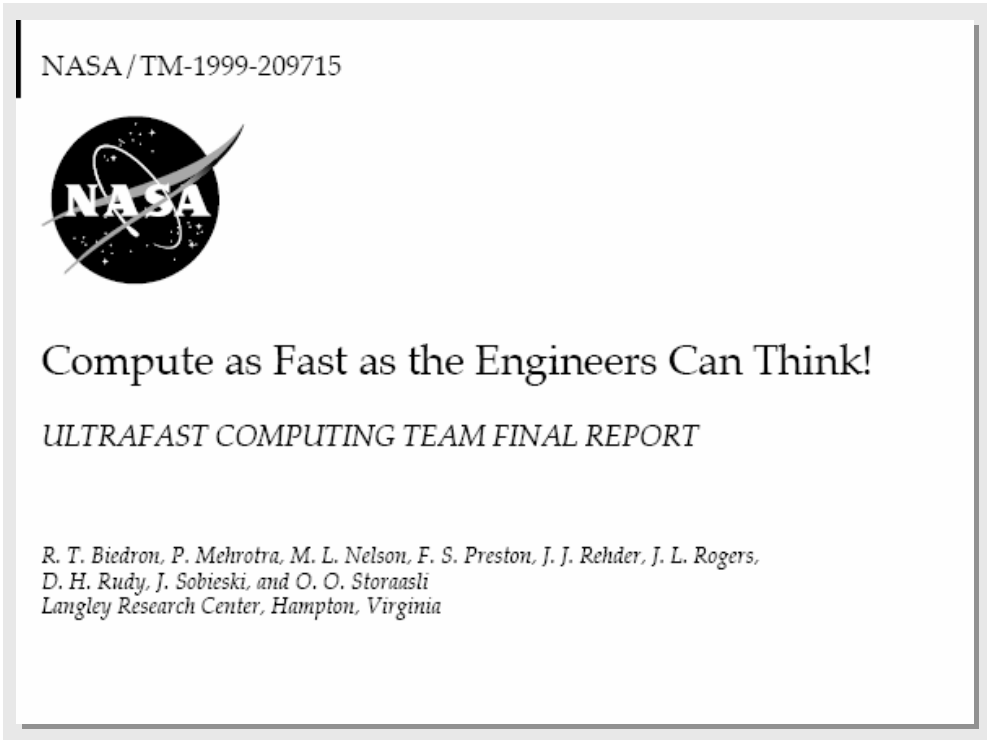

NASA/TM-1999-209715

Compute as Fast as the Engineers Can Think!

ULTRAFAST COMPUTING TEAM FINAL REPORT

R. T. Biedron, P. Mehrotra, M. L. Nelson, F. S. Preston, J. J. Rehder, J. L. Rogers, D. H. Rudy, J. Sobieski, and O. O. Storaasli
Langley Research Center, Hampton, Virginia

Sandia National Laboratories

Outline

- Applications of the Future
- Limits of Moore's Law
- An Expert System/Optimizer for Supercomputing
- Reaching to Zettaflops
- Roadmap and Future Directions

Sandia National Laboratories

*** This is a Preview ***

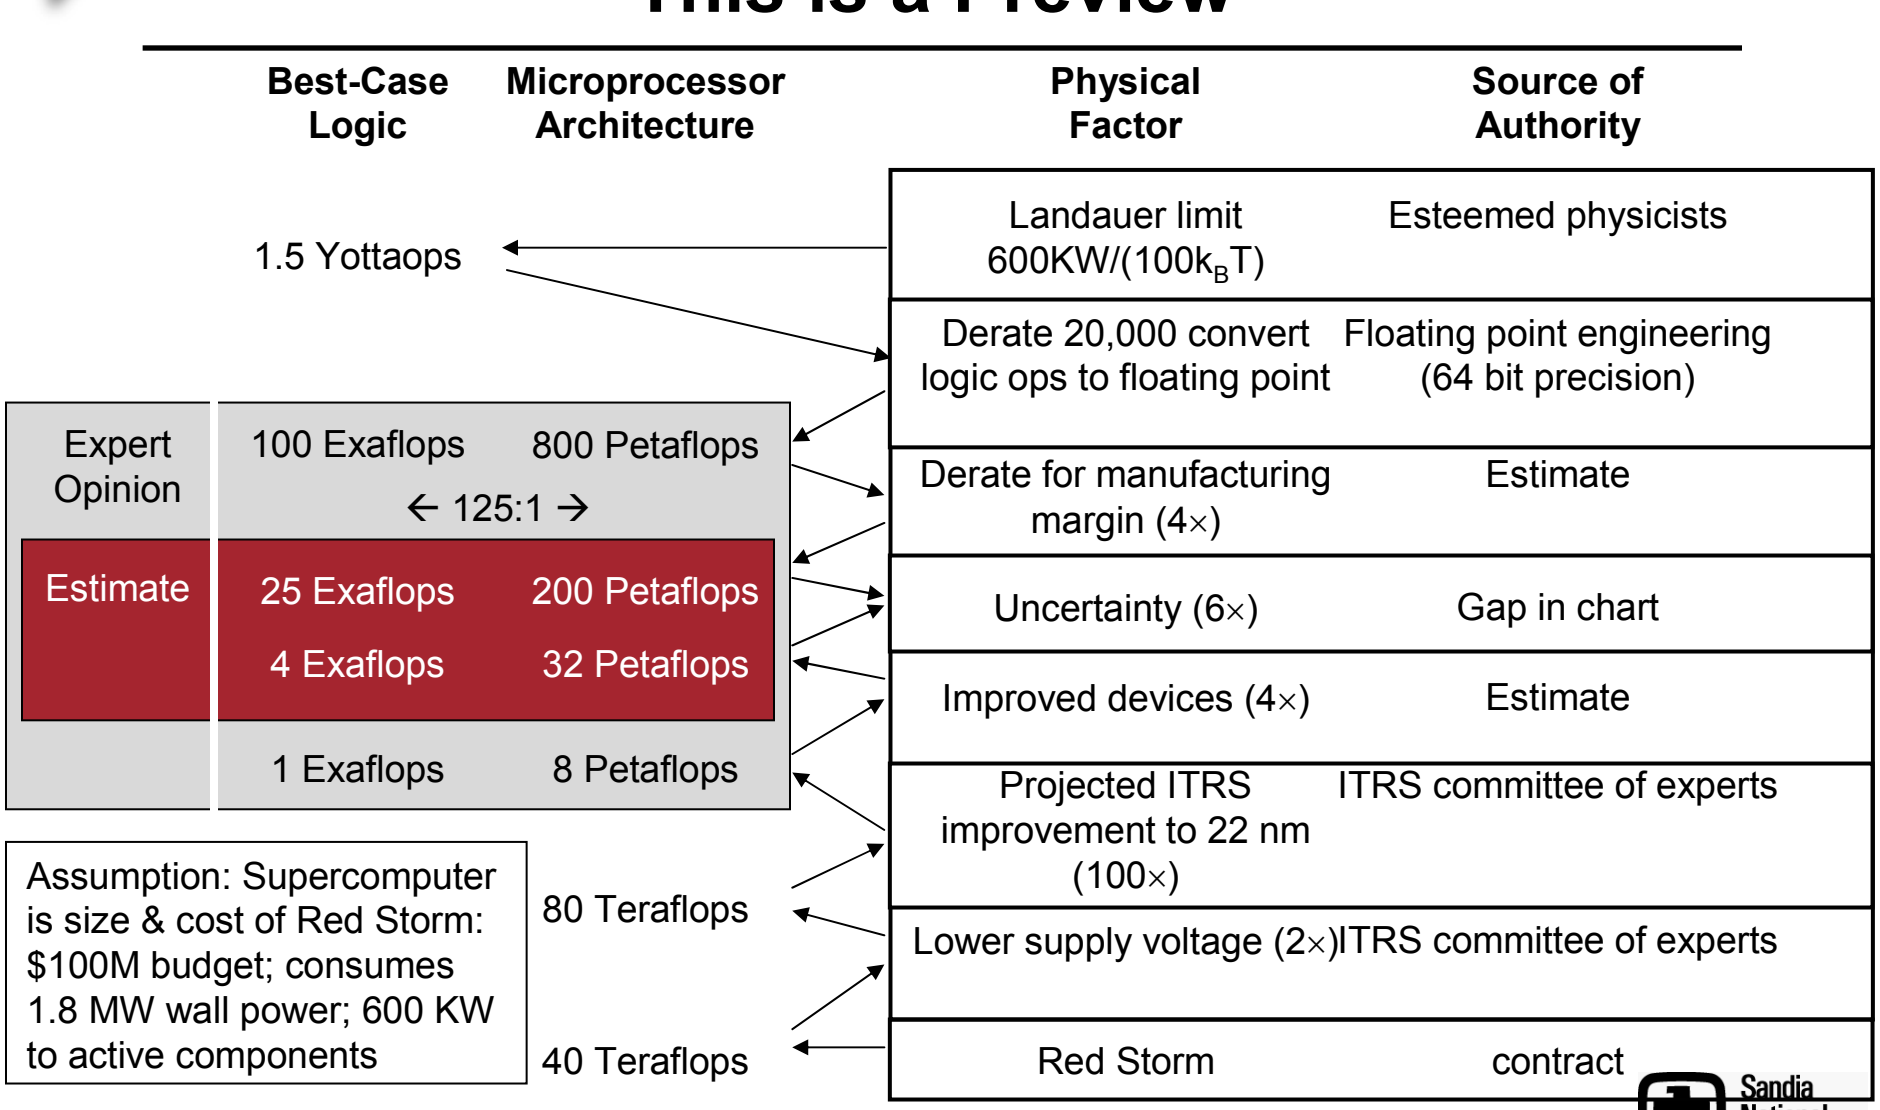

Sandia National Laboratories

Thermal Noise Limit

This logical irreversibility is associated with physical irreversibility and requires a minimal heat generation, per machine cycle, typically of the order of kT for each irreversible function.
– R. Landauer 1961

Irreversibility and Heat Generation in the Computing Process

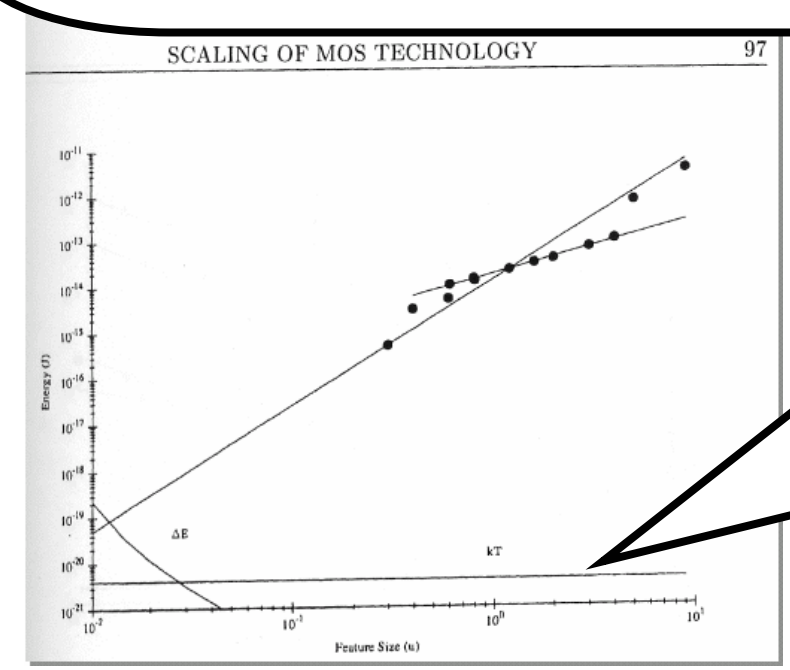

kT "helper line," drawn out of the reader's focus because it wasn't important at the time of writing
– Carver Mead, Scaling of MOS Technology, 1994

Sandia National Laboratories

Metaphor: FM Radio on Trip to Sacramento

- You drive to Sacramento listening to FM radio
- Music clear for a while, but noise creeps in and then overtakes music
- Why?
  - Signal at antenna weakens
  - Thermal electron noise constant at k_B T
- Analogy: You live out the next dozen years buying PCs every couple years
- Electrical effect
  - Moore's Law causes switching energy of gates to decrease at about 30% per year
  - Thermal electron noise constant at k_B T

Details: Erik DeBenedictis, "Taking ASCI Supercomputing to the End Game," SAND2004-0959

Sandia National Laboratories

FM Radio and End of Moore's Law

Distance

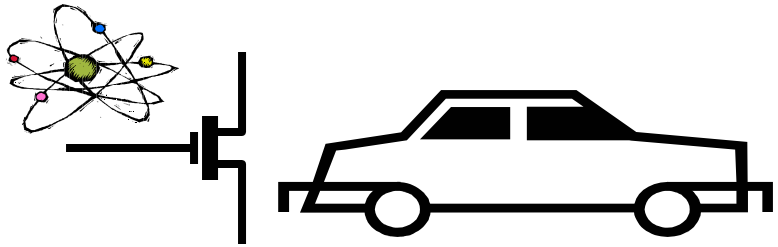

Driving away from FM transmitter→less signal
Noise from electrons → no change

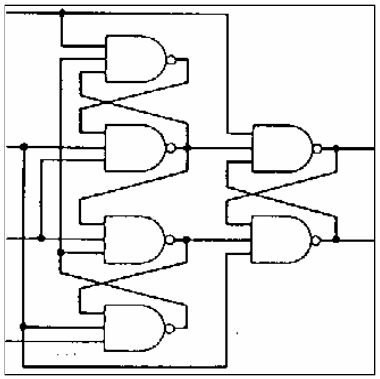

Shrink

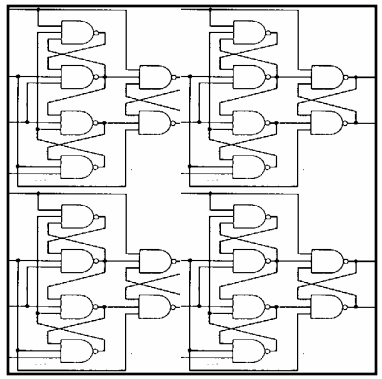

Increasing numbers of gates→less signal power
Noise from electrons → no change

Sandia National Laboratories

Amount of Reliability Needed

- We expect computers to be reliable
- A future supercomputer will perform 10^{30}-10^{40} operations in its lifetime
- Error rate should be < 10^{-30} - 10^{-40}
- Reliability due to thermal noise about $e^{-E/kt}$
- Need about e^{-100} error rate, or 100 k_BT switching energy

SNR (db)	Power Ratio	P_{error}
10	10	3.9×10^{-6}
14	25	6.8×10^{-13}
18	63	1.4×10^{-29}
22	160	3.3×10^{-71}
26	400	1.8×10^{-175}
30	1,000	4.5×10^{-437}
34	2,500	7.1×10^{-1094}
38	6,300	2.2×10^{-2743}
42	16,000	1.8×10^{-6886}
46	40,000	3.8×10^{-17293}
50	100,000	3.2×10^{-43433}
54	250,000	8.1×10^{-10194}
58	630,000	1.8×10^{-274025}
62	1,500,000	9.6×10^{-688315}

Noise Limit

2016

Today

$$q := \int_t^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx;\ t \to \sqrt{2 * 10^{\frac{snr}{10}}}$$

Sandia National Laboratories

Noise Levels

- 0 db Limit of hearing
- 20 db Rustling leaves
- 40-50 db Typical neighborhood
- 60-70 db Normal conversation
- 80 db Telephone dial tone
- 85 db City traffic inside car
- 90 db Train whistle @500’
- 95 db Subway train @200’
- 90-95 db Ear damage

- Today: 50 db
  - Thermal noise:Logic:: Rustling leaves:Talking
- 2016: 30 db
  - Thermal noise:Logic:: Talking:Train Whistle
- Reliability limit 20 db
  - Thermal noise:Logic:: Outside neighborhood:Talking

Personal Observational Evidence

- Have radios become better able to receive distant stations over the last few decades with a rate of improvement similar to Moore's Law?
- You judge from your experience, but the answer should be that they have not.
- Therefore, electrical noise does not scale with Moore's Law.

Sandia National Laboratories

SIA Semiconductor Roadmap

- Generalization of Moore's Law
  - Projects many parameters
  - Years through 2016
  - Includes justification
  - Panel of experts
    - known to be wrong
  - Size between Albuquerque white and yellow pages

International Technology Roadmap for Semiconductors (ITRS), see http://public.itrs.net

Sandia National Laboratories

Semiconductor Roadmap

YEAR OF PRODUCTION	2010	2013	2016
DRAM ½ PITCH (nm)	45	32	22
MPU / ASIC ½ PITCH (nm)	50	35	25
MPU PRINTED GATE LENGTH (nm)	25	18	13
MPU PHYSICAL GATE LENGTH (nm)	18	13	9
Physical gate length high-performance (HP) (nm) [1]	18	13	9
Equivalent physical oxide thickness for high-performance T_{ox} (EOT)(nm) [2]	0.5-0.8	0.4-0.6	0.4-0.5
Gate depletion and quantum effects electrical thickness adjustment factor (nm) [3]	0.5	0.5	0.5
T_{ox} electrical equivalent (nm) [4]	1.2	1.0	0.9
Nominal power supply voltage (V_{dd}) (V) [5]	0.6	0.5	0.4
Nominal high-performance NMOS sub threshold leakage current, $I_{sd,leak}$ (at 25°C) (μA/μm) [6]	3	7	10
Nominal high-performance NMOS saturation drive current , I_{dd} (at V_{dd}, at 25°C) (μA/μm) [7]	1200	1500	1500
Required percent current-drive "mobility/transconductance improvement" [8]	30%	70%	100%
Parasitic source/drain resistance (Rsd) (ohm-μm) [9]	110	90	80
Parasitic source/drain resistance (Rsd) pe	25%	30%	35%
Parasitic capacitance percent of ideal gat	31%	36%	42%
High-performance NMOS device τ (C_{gate} * V_{dd} / I_{dd}-NMOS)(ps) [12]	0.39	0.22	0.15
Relative device performance [13]	4.5	7.2	10.7
Energy per (W/L_{gate}=3) device switching transition (C_{gate}*(3*L_{gate})*V^2) (fJ/Device) [14]	0.015	0.007	0.002
Static power dissipation per (W/Lgate=3) device (Watts/Device) [15]	9.7E-08	1.4E-07	1.1E-07

1,000 k_BT/transistor

White—Manufacturable Solutions Exist, and Are Being Optimized

Yellow—Manufacturable Solutions are Known

Red—Manufacturable Solutions are NOT Known

Sandia National Laboratories

Limits for a $100M-Class Supercomputer

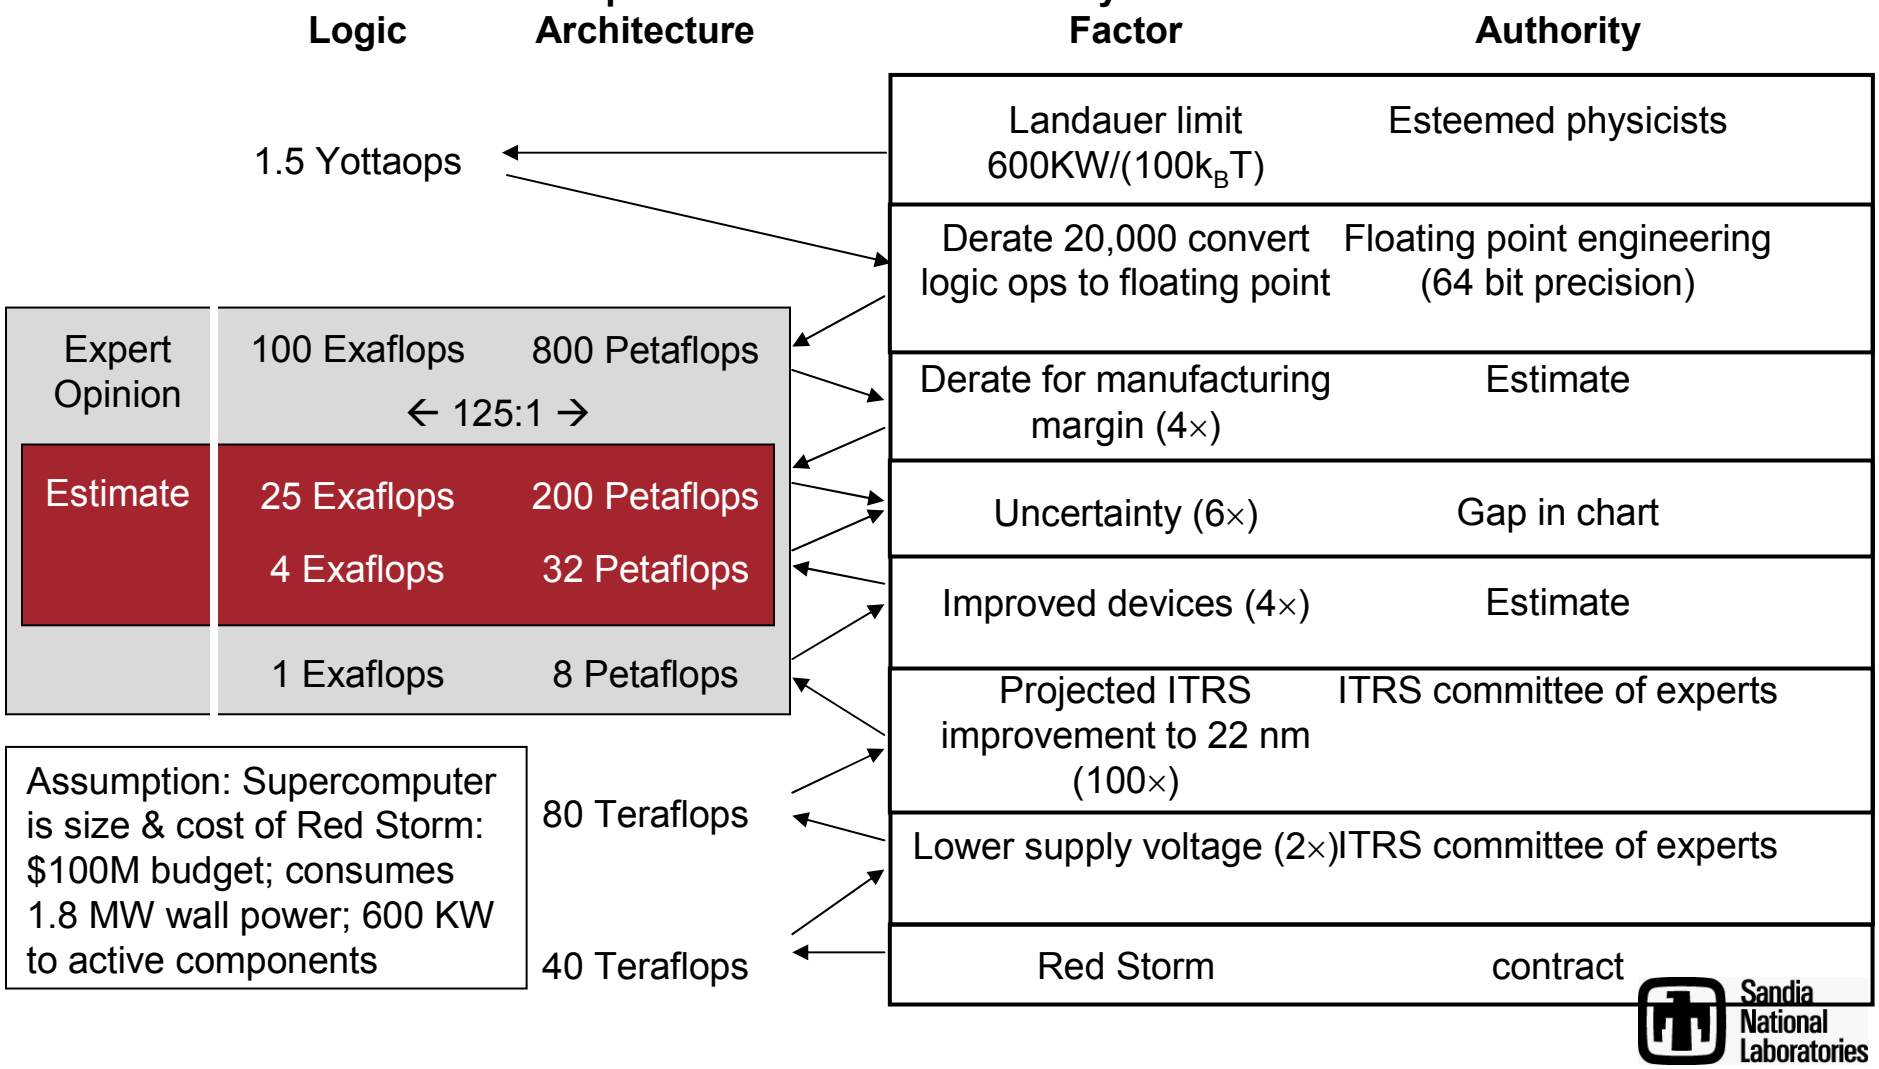

Sandia National Laboratories

Outline

- Applications of the Future
- Limits of Moore's Law
- An Expert System/Optimizer for Supercomputing
- Reaching to Zettaflops
- Roadmap and Future Directions

Sandia National Laboratories

Expert System for Future Supercomputers

- Applications Modeling
  - Runtime
    $T_{run} = f_1(n, design)$
- Technology Roadmap
  - Gate speed = $f_2(year)$,
  - chip density = $f_3(year)$,
  - cost = $(n, design), …
- Scaling Objective Function
  - I have $\$C_1$ & can wait $T_{run}=C_2$ seconds. What is the biggest n I can solve in year Y?

- Use "Expert System" To Calculate:

$$\underset{\text{All designs}}{\text{Max}} \quad n: \$<C_1,\ T_{run}<C_2$$

- Report:

$$\frac{\text{Floating operations}}{T_{run}(n, design)}$$

and illustrate "design"

Sandia National Laboratories

Analytical Runtime Model

- Simple case: finite difference equation
- Each node holds n×n×n grid points

- Volume-area rule
  - Computing $\propto n^3$
  - Communications $\propto n^2$

$$T_{step} = 6\, n^2\, C_{bytes}\, T_{byte} + n^3\, F_{grind}/floprate$$

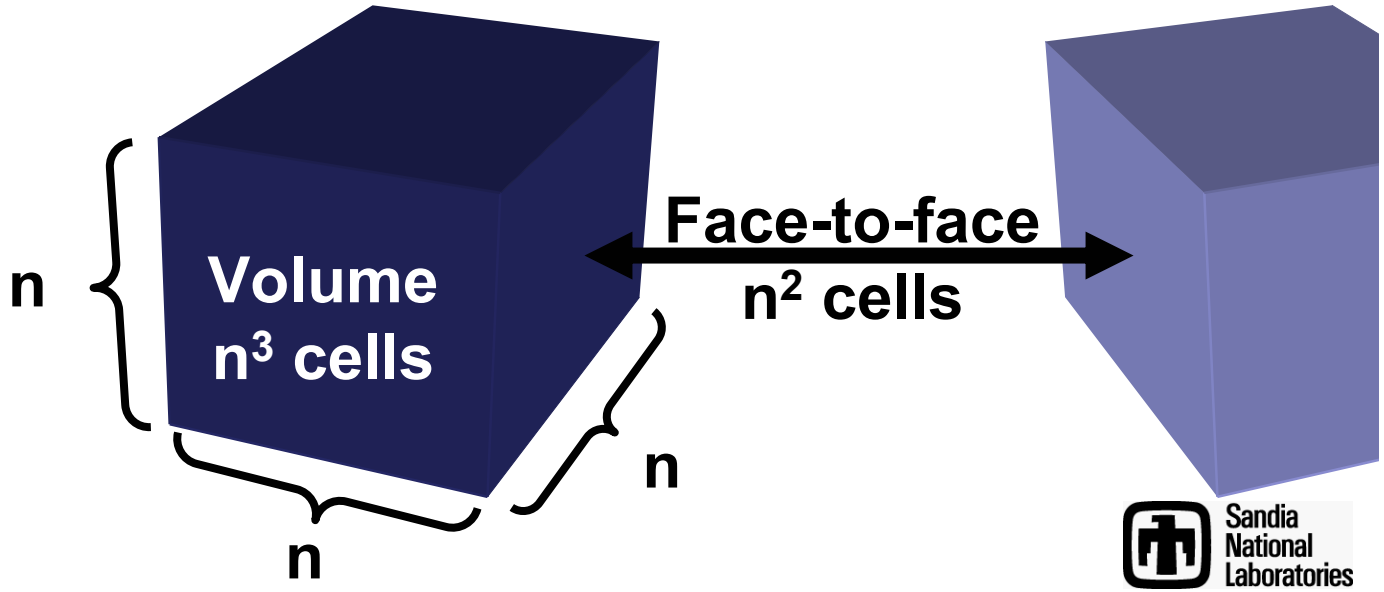

Sandia National Laboratories

Supercomputer Expert System

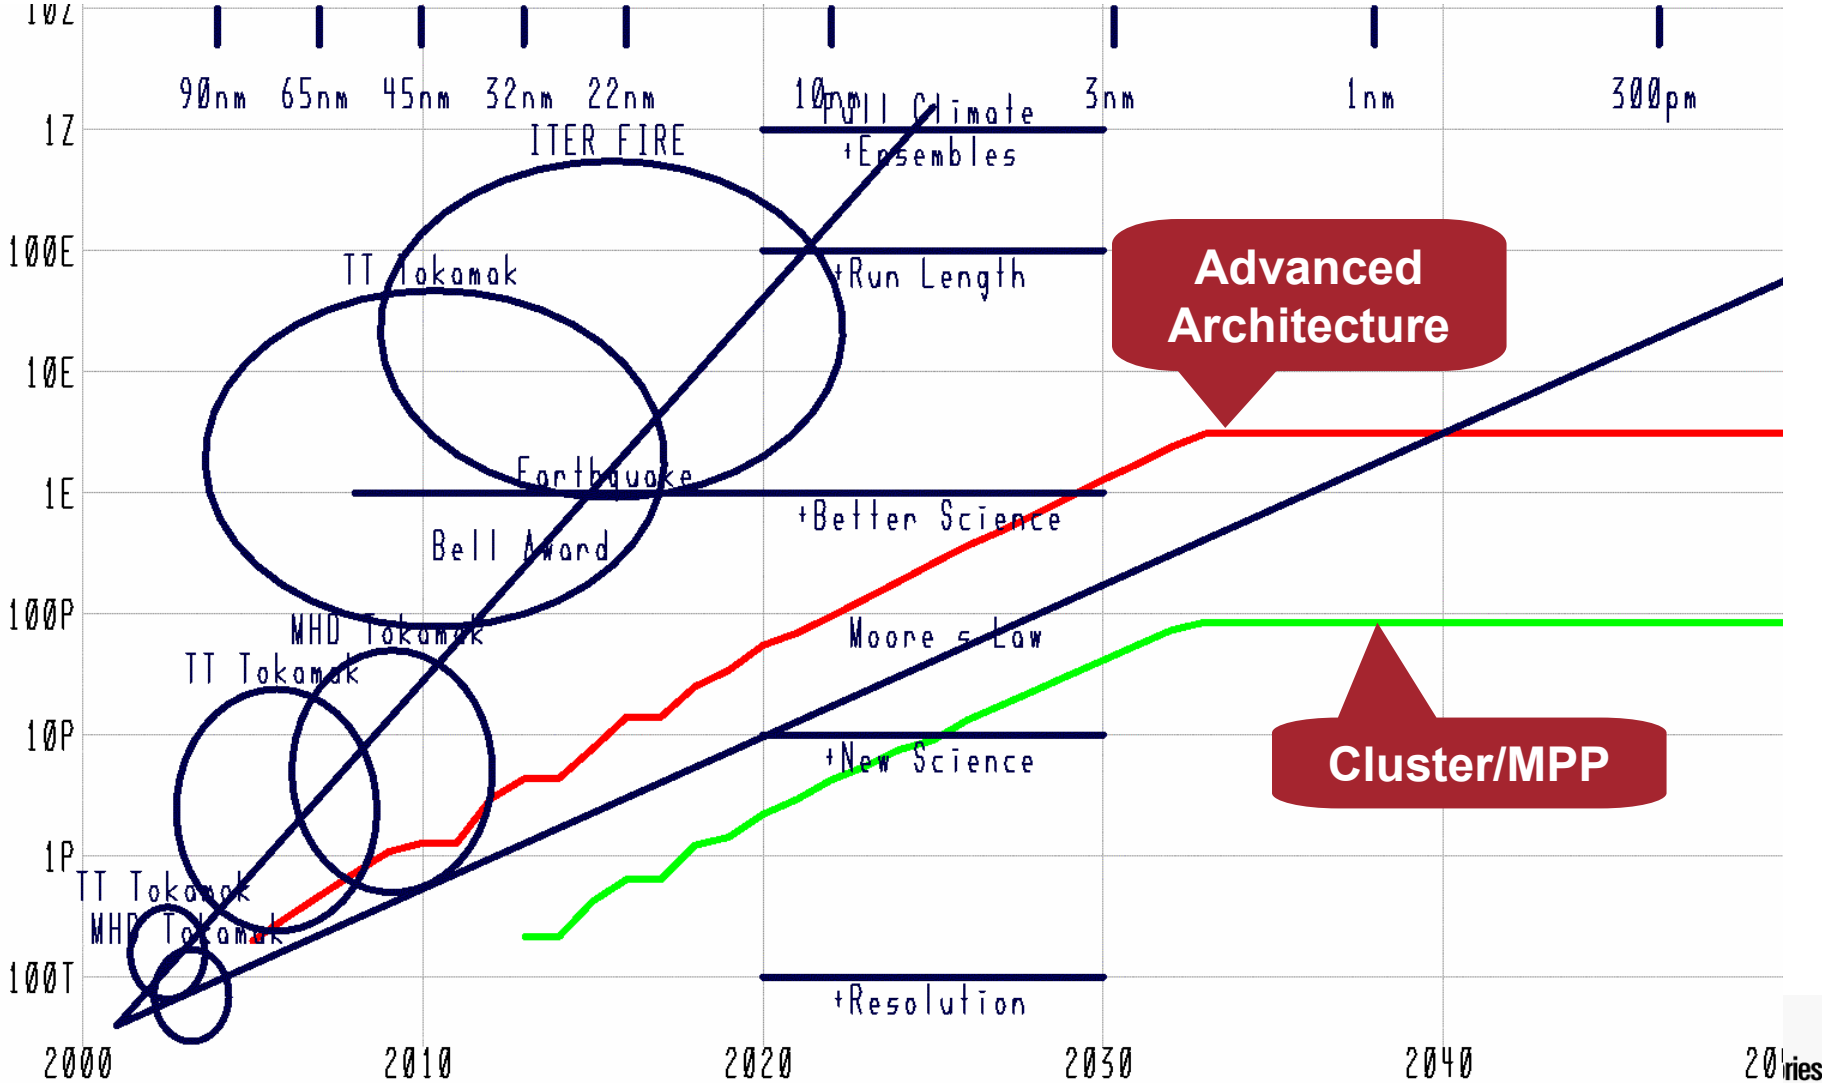

Outline

- Applications of the Future
- Limits of Moore's Law
- An Expert System/Optimizer for Supercomputing
- Reaching to Zettaflops
- Roadmap and Future Directions

Sandia National Laboratories

*** This is a Preview ***

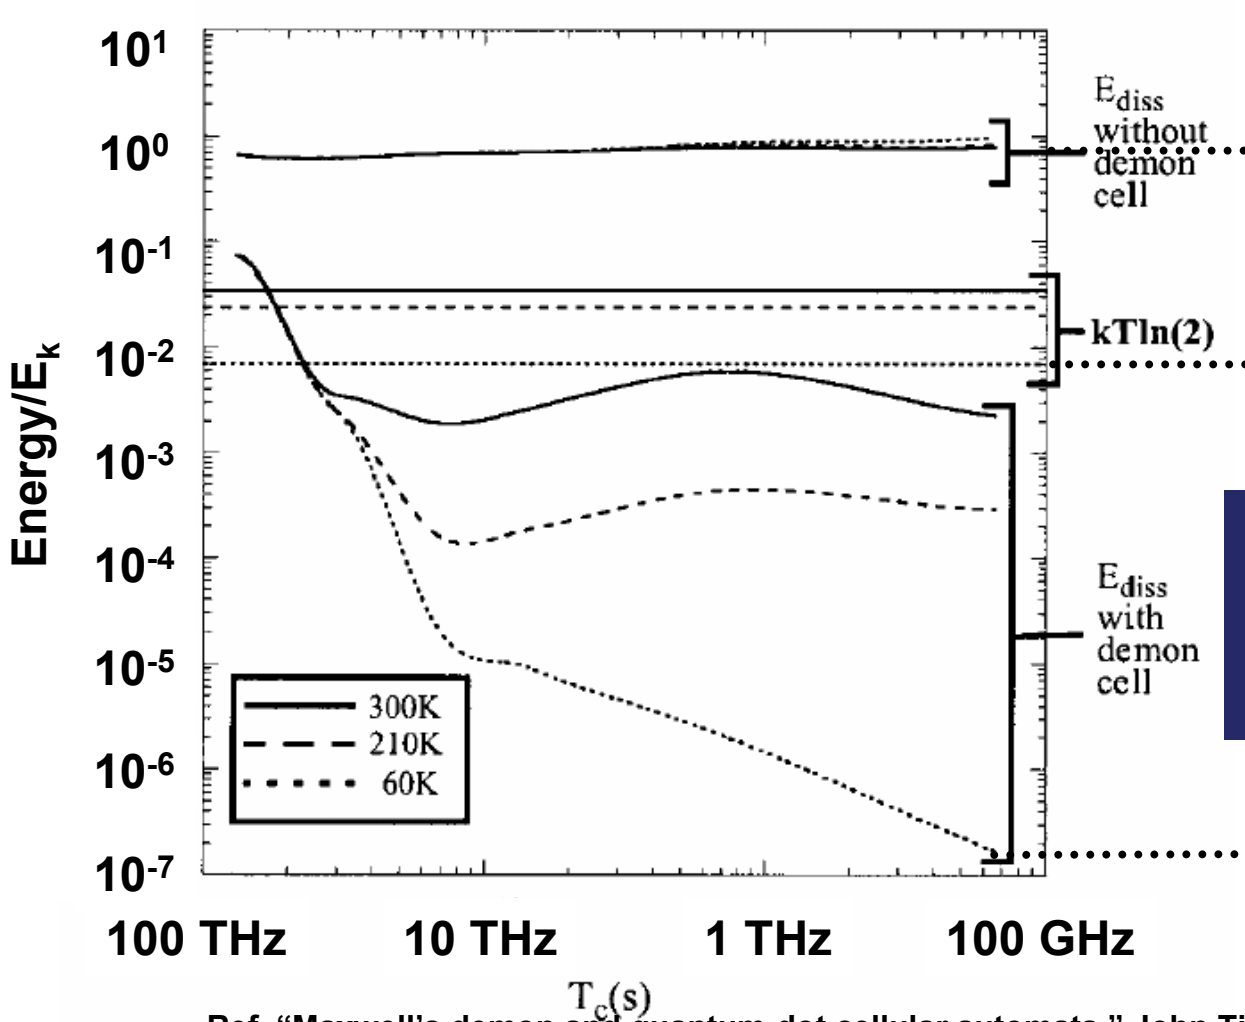

2004 Device Level

1000 ×

"Reliability Limit"

150 ×

"Landauer Limit"

>10^4 × Improvement @ 100 GHz & 60° K

Dissipation for reversible operations

Ref. "Maxwell's demon and quantum-dot cellular automata," John Timler and Craig S. Lent, JOURNAL OF APPLIED PHYSICS 15 JULY 2003

Sandia National Laboratories

Today's Universal Logic & Reliability Limit

- Today's logic operates on a simple principle
  - Create a "1" by taking charge from the positive supply
  - Create a "0" by sending charge to the negative supply
- Energy Consumption
  - Each gate switch generates $E_{sw} = \frac{1}{2} CV^2 > \sim 100 k_B T$ heat

Signal energy must be greater than ~100 $k_B T$ to avoid spontaneous glitches. To change a bit, convert energy to heat.

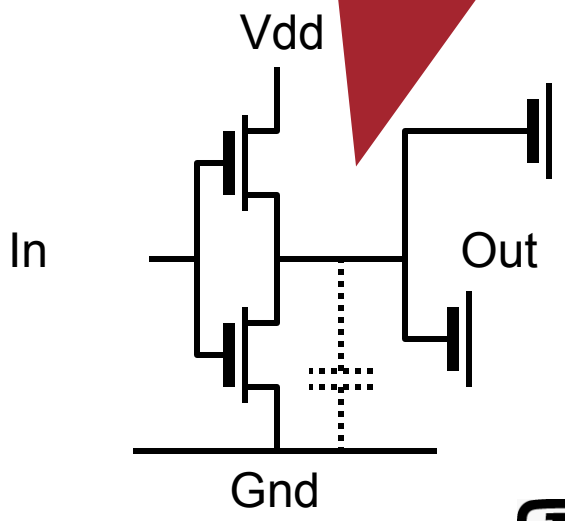

Sandia National Laboratories

“Recycling” Power

- The 100k_BT limit appears unbeatable, but the energy can be “recycled”
- Diagram shows a “SCRL” circuit with regular transistors
- Power comes through a largely loss less resonant device (tuning fork)
- No apology offered for the mechanical device; this is the price of progress

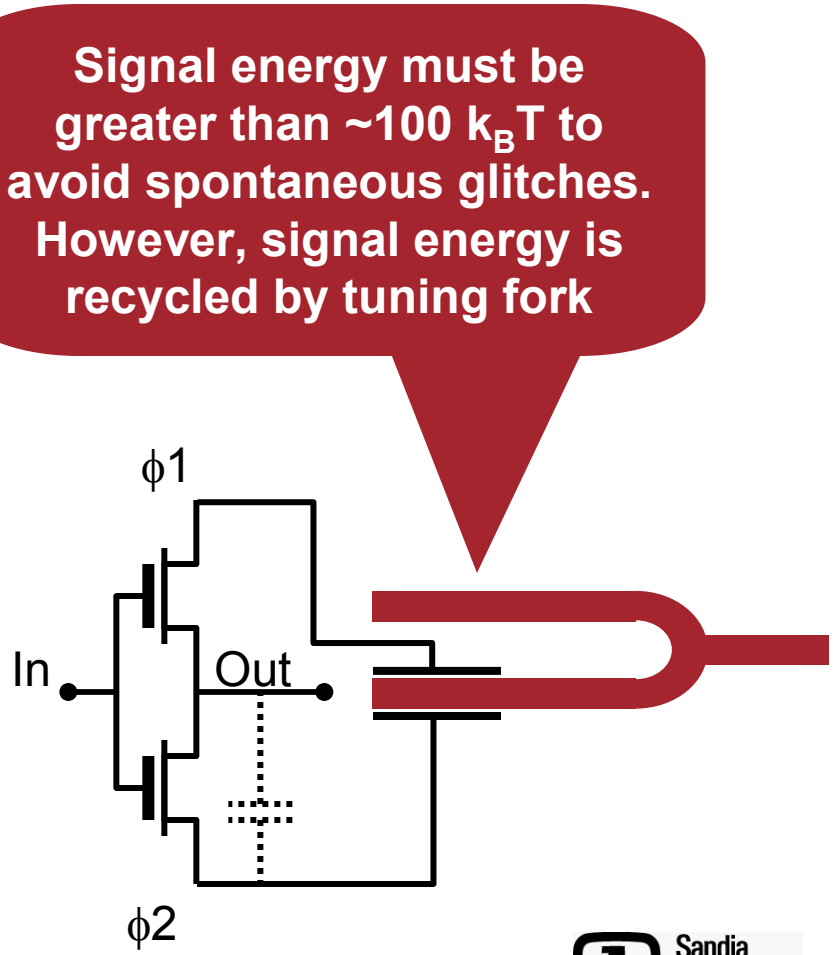

Sandia National Laboratories

Resonant Clocks

- **Tuning Fork**
  - **Nice idea but slow**
- **MEMs Resonator**
  - **Moderate speed and compatible with silicon fabrication**

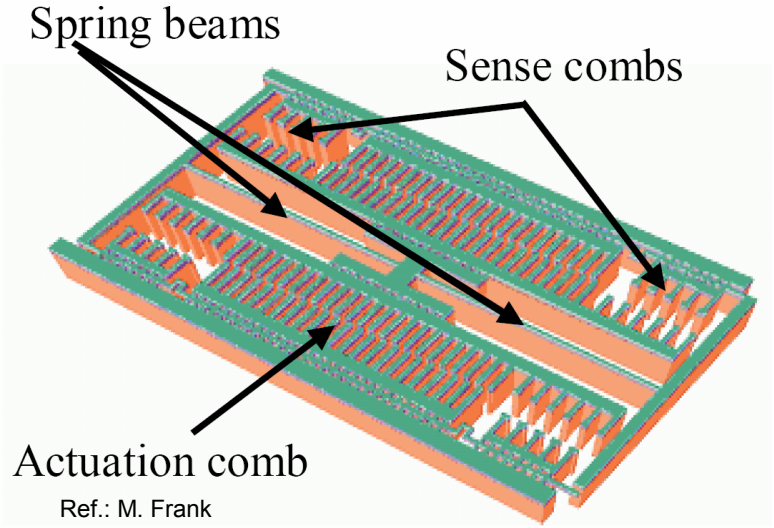

Ref.: M. Frank

- **Carbon Nanotube**
  - **Simulated to 50 GHz but not known how to fabricate at present**

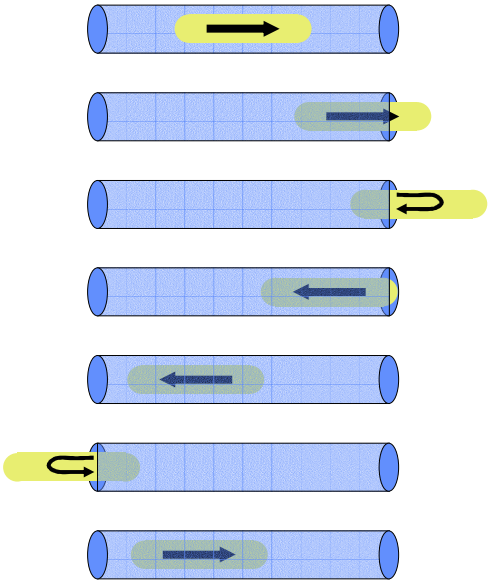

Sandia National Laboratories

Landauer's Arguments

- **Landauer makes three arguments in his 1961 paper**
  - **Kintetics of a bistable well (next slide)**
  - **Entropy generation →**

Sorry, I don't have a cute story (like the FM radio) for Landauer's argument

- **Entropy of a system in statistical mechanics:**

$$S = k_B \log_e(W)$$

W is number of states

- **Entropy of a mechanical system containing a flip flop in an unknown state:**

$$S = k_B \log_e(2W)$$

- **After clearing the flip flop:**

$$S = k_B \log_e(W)$$

- **Difference $k_B \log_e(2)$**

Sandia National Laboratories

Landauer's Limit

- The Landauer limit says you can reduce power dissipation for irreversible functions below 100 $k_B T$, but not below $k_B T \log_e 2$
- In the diagram on the right, when the energy barrier drops to below about $k_B T$, the state will spontaneously switch and dissipate remaining energy as heat

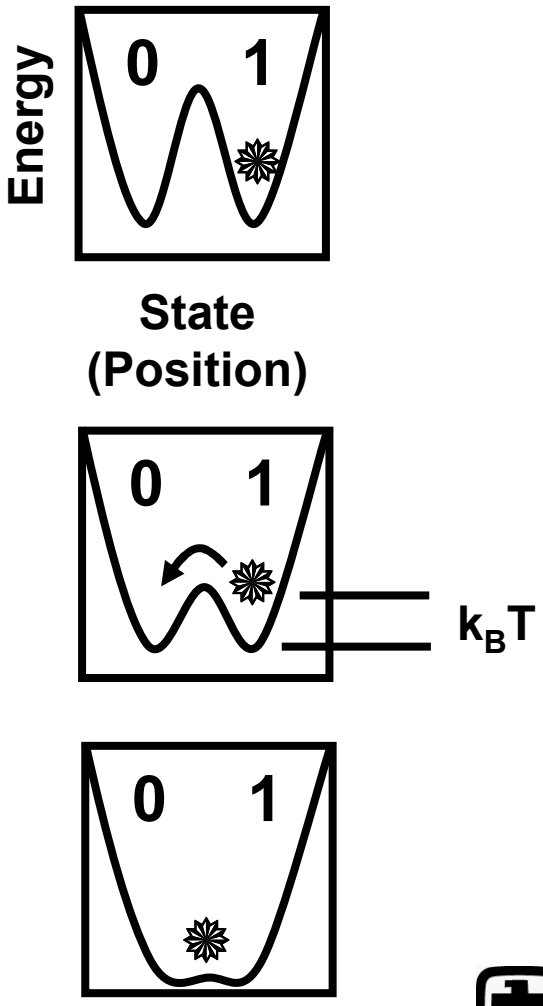

Sandia National Laboratories

Reversible Logic Can Beat Landauer's Limit

- Any function can be made reversible by saving its inputs
- Diagram below outlines an asymptotically zero-energy way to perform the AND function, in composition with other logical operations

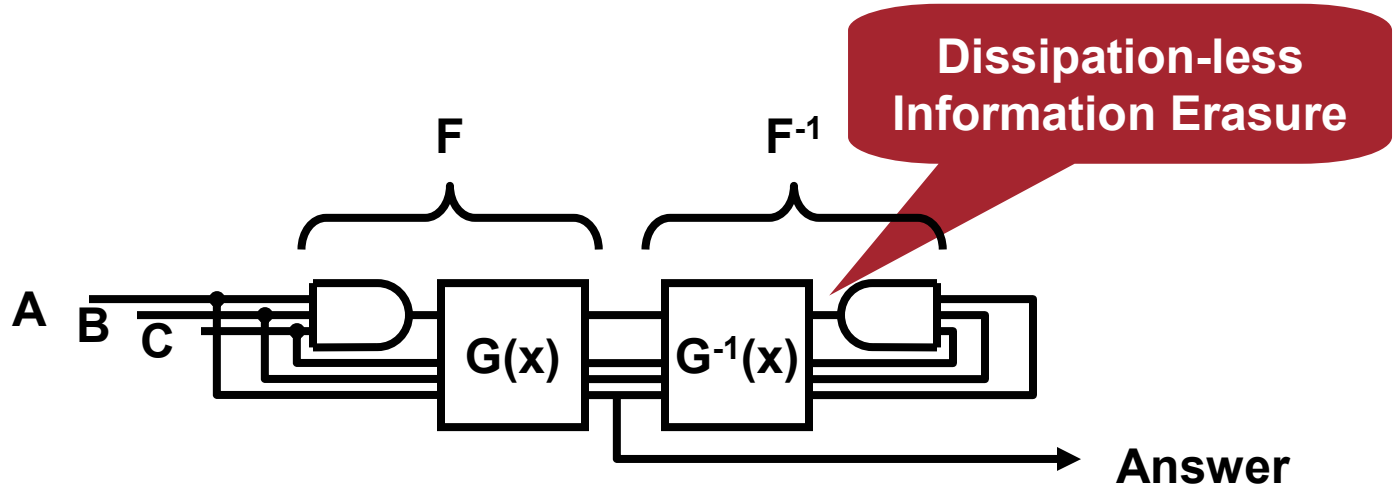

Sandia National Laboratories

A New Computing Device: Quantum Dots

- Pairs of molecules create a memory cell or a logic gate

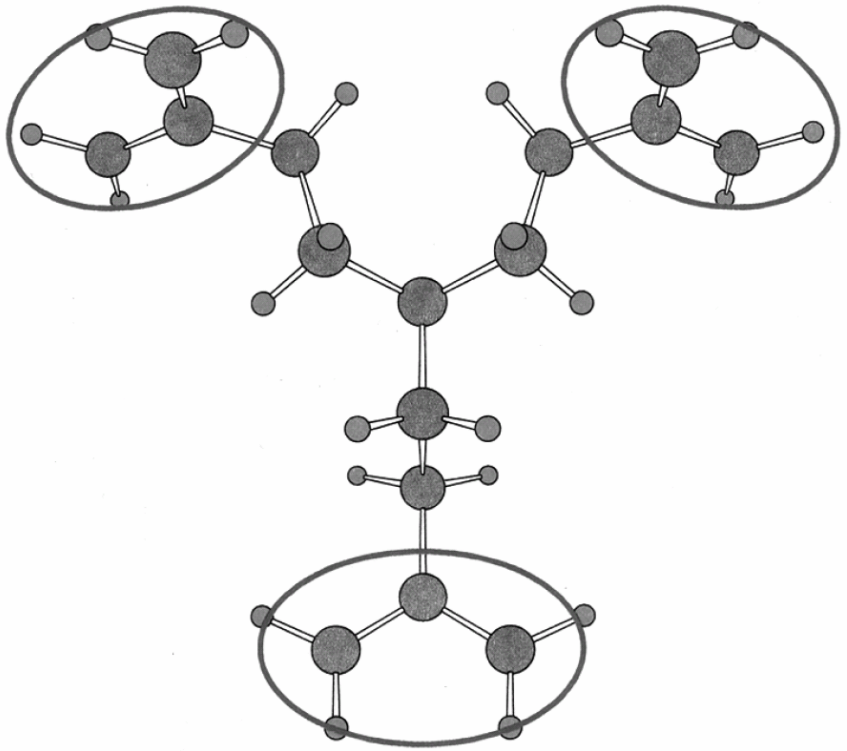

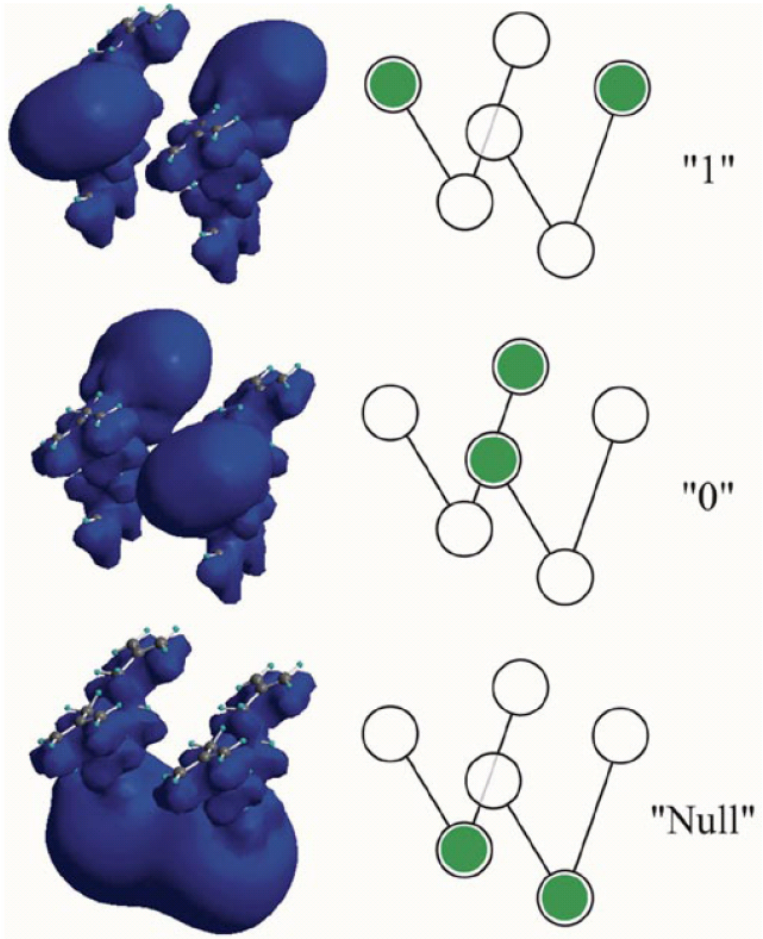

Ref. "Clocked Molecular Quantum-Dot Cellular Automata," Craig S. Lent and Beth Isaksen, IEEE TRANSACTIONS ON ELECTRON DEVICES, VOL. 50, NO. 9, SEPTEMBER 2003

Sandia National Laboratories

Upside Potential of Quantum Dots

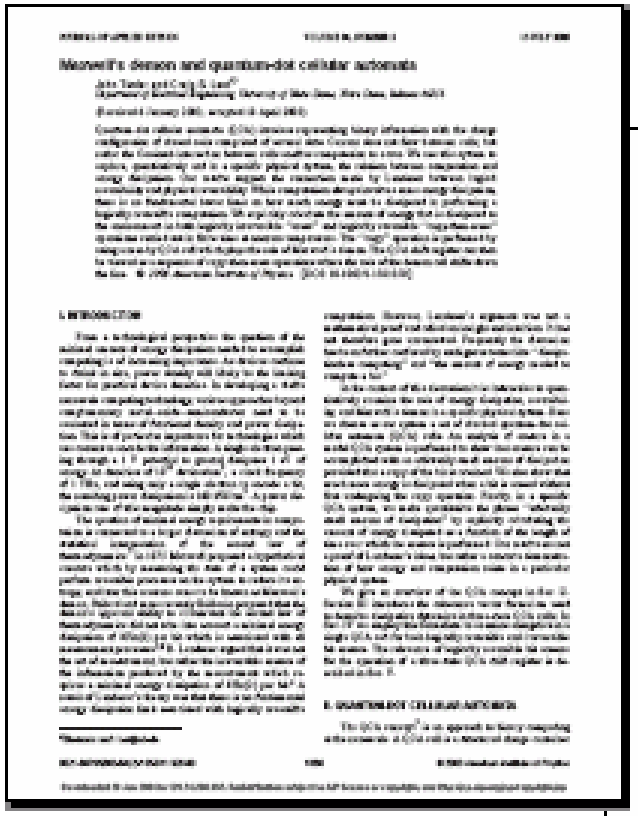

Maxwell's demon and quantum-dot cellular automata

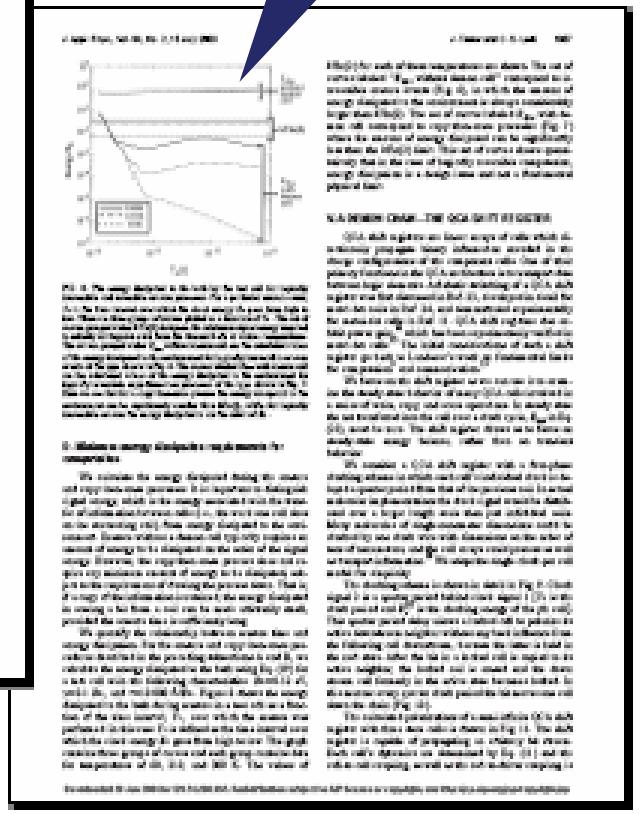

Next Slide

Ref. "Maxwell's demon and quantum-dot cellular automata," John Timler and Craig S. Lent, JOURNAL OF APPLIED PHYSICS 15 JULY 2003

Sandia National Laboratories

Upside Potential of Quantum Dots

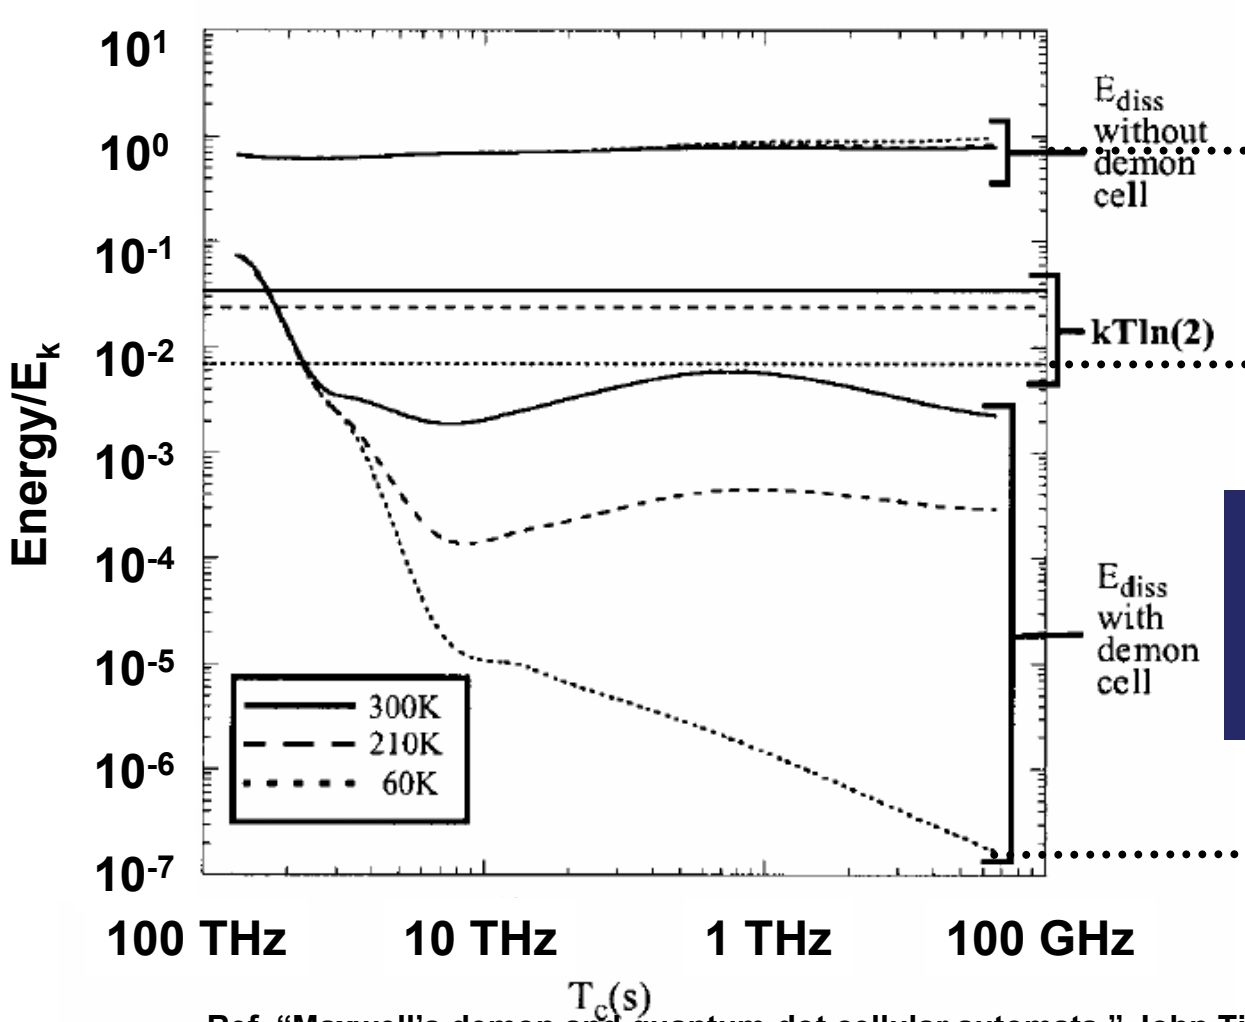

2004 Device Level

1000 ×

"Reliability Limit"

150 ×

"Landauer Limit"

>10^4 × Improvement @ 100 GHz & 60° K

Dissipation for reversible operations

Ref. "Maxwell's demon and quantum-dot cellular automata," John Timler and Craig S. Lent, JOURNAL OF APPLIED PHYSICS 15 JULY 2003

Sandia National Laboratories

Reversible Multiplier Status

- 8×8 Multiplier Designed, Fabricated, and Tested by IBM & University of Michigan
- Power savings was up to 4:1

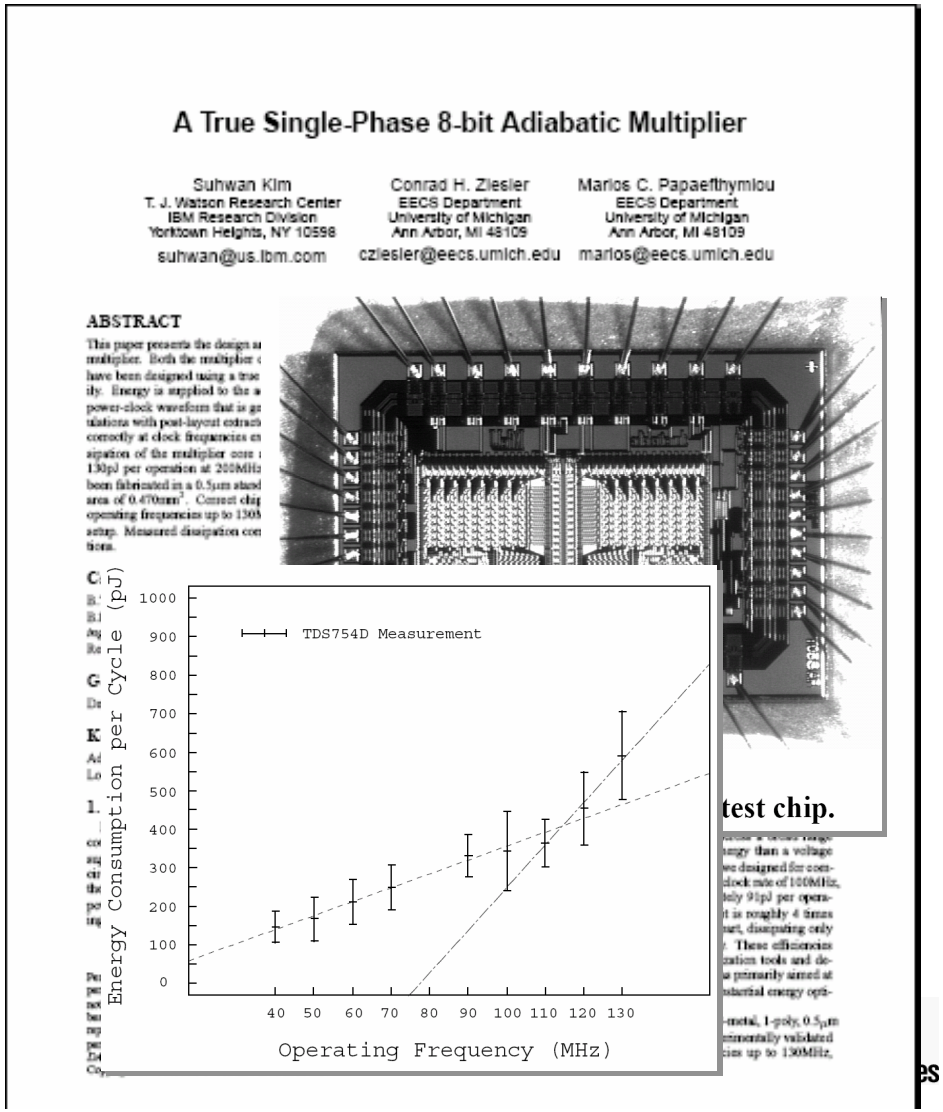

A True Single-Phase 8-bit Adiabatic Multiplier

Suhwan Kim
T. J. Watson Research Center
IBM Research Division
Yorktown Heights, NY 10598
suhwan@us.ibm.com

Conrad H. Ziesler
EECS Department
University of Michigan
Ann Arbor, MI 48109
cziesler@eecs.umich.edu

Marios C. Papaefthymiou
EECS Department
University of Michigan
Ann Arbor, MI 48109
marios@eecs.umich.edu

ABSTRACT

test chip.

Reversible Microprocessor Status

- Status
  - Subject of Ph. D. thesis
  - Chip laid out (no floating point)
  - RISC instruction set
  - C-like language
  - Compiler
  - Demonstrated on a PDE
  - However: really weird and not general to program with +=, -=, etc. rather than =

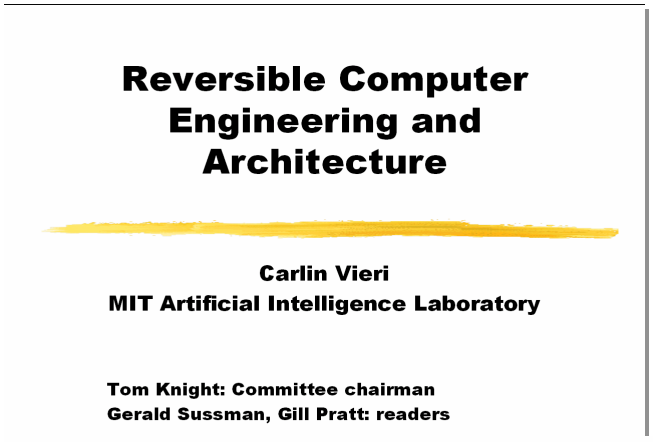

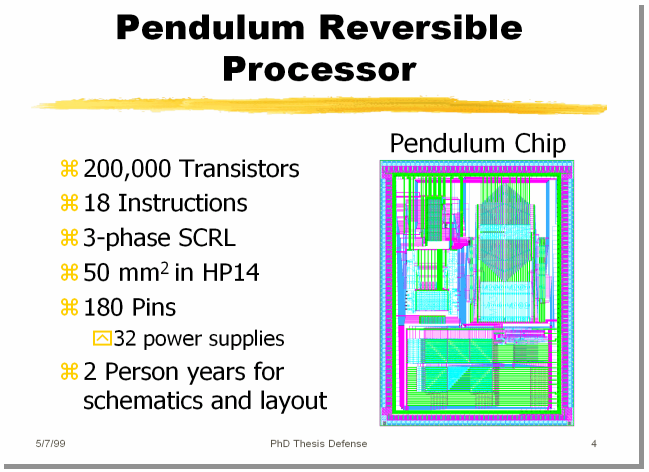

Sandia National Laboratories

Supercomputer Expert System

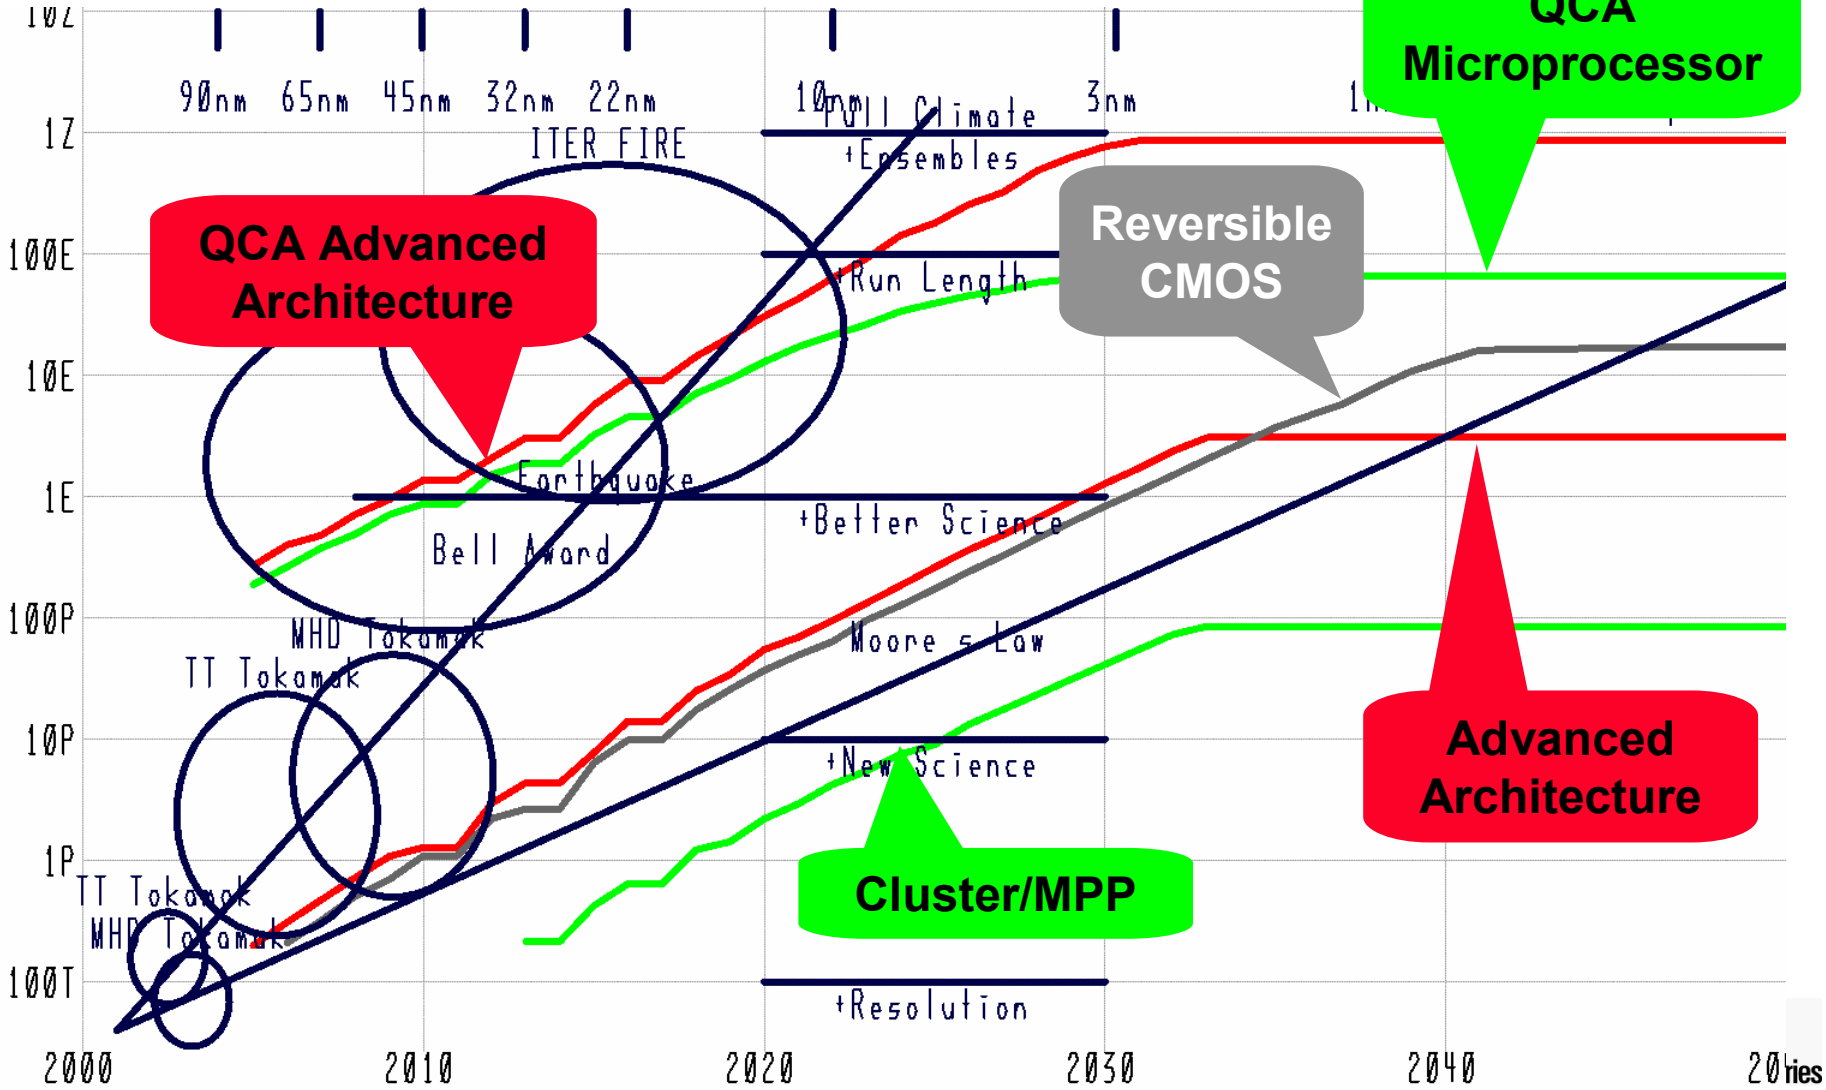

Other Approaches to Extreme Computing

- Other fairly normal devices
  - Ballistic Y junction, parametric quantron
- Nanotechnology
  - New technologies are easy to propose when you can do atom-by-atom assembly. Helical logic, Carbon nanotube Y junctions, rod logic, …
- Use the third dimension
  - Chips are 1cm x 1cm x 2 nm on .2 cm centers
  - Increase the 2 nm dimension!
- Coherent quantum computing

Sandia National Laboratories

Outline

- Applications of the Future
- Limits of Moore's Law
- An Expert System/Optimizer for Supercomputing
- Reaching to Zettaflops
- Roadmap and Future Directions

Sandia National Laboratories

Conclusions

- A review of supercomputer applications reveals a continuum up to 1 Zettaflops
- A review of "the physics of computation" reveals a progression of technologies offering progressively more performance for progressively more development effort
  - with no theoretical upper bound
- Integrating the ideas (with an expert system) suggests they could work together in a supercomputer
  - This is not the same as proving the ideas right!

Sandia National Laboratories

Where to Go Next I: Workshop

- Don’t believe me? Believe the Experts
- Workshop Agenda October
  - Applications session – Climate expert Phil Jones
  - Advanced Architectures – PIM expert Peter Kogge
  - Limits of Current Architectures – Me
  - Limits Panel I: Limits of Current Technology
  - New Logic – Reversible Logic Expert Michael Frank
  - New Devices – Quantum Dot Developer Craig Lent
  - Limits Panel II: Opportunities with Innovation

Where To Go Next II: Roadmap

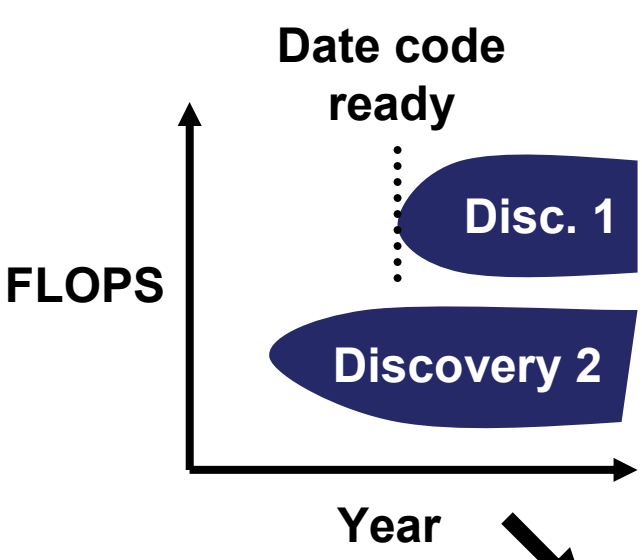

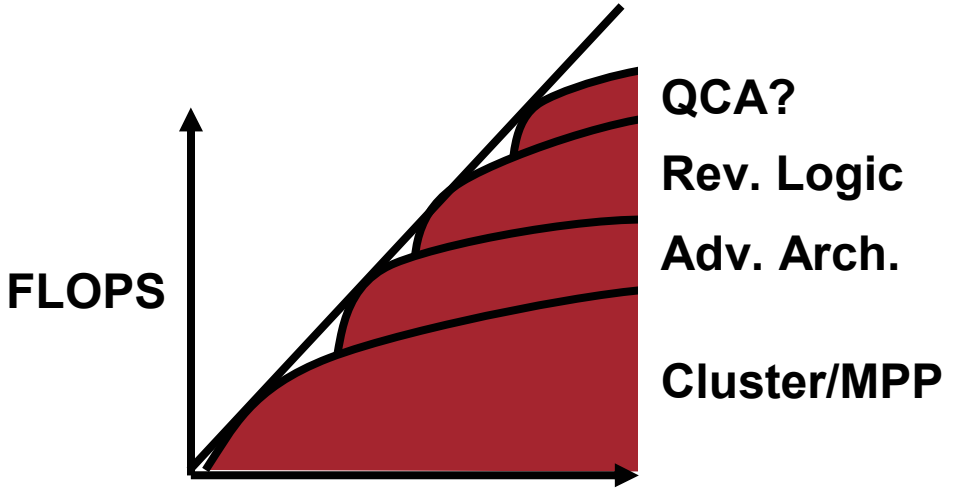

Notes:

* Not necessarily one machine; different applications may require different machines
* Specifics are just my ideas

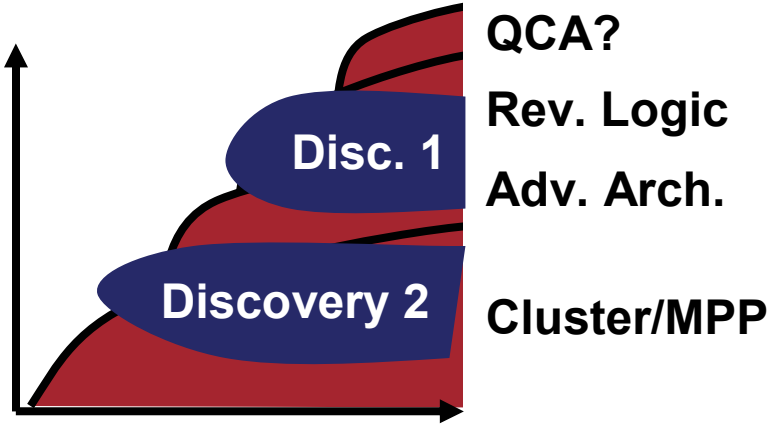

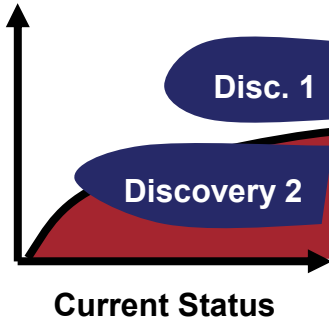

Current Status

Sandia National Laboratories